

Learning probabilities from random observables in high dimensions: the maximum entropy distribution and others

Tomoyuki Obuchi · Simona Cocco · Rémi Monasson

the date of receipt and acceptance should be inserted later

Abstract We consider the problem of learning a target probability distribution over a set of N binary variables from the knowledge of the expectation values (with this target distribution) of M observables, drawn uniformly at random. The space of all probability distributions compatible with these M expectation values within some fixed accuracy, called version space, is studied. We introduce a biased measure over the version space, which gives a boost increasing exponentially with the entropy of the distributions and with an arbitrary inverse ‘temperature’ Γ . The choice of Γ allows us to interpolate smoothly between the unbiased measure over all distributions in the version space ($\Gamma = 0$) and the pointwise measure concentrated at the maximum entropy distribution ($\Gamma \rightarrow \infty$). Using the replica method we compute the volume of the version space and other quantities of interest, such as the distance R between the target distribution and the center-of-mass distribution over the version space, as functions of $\alpha = (\log M)/N$ and Γ for large N . Phase transitions at critical values of α are found, corresponding to qualitative improvements in the learning of the target distribution and to the decrease of the distance R . However, for fixed α , the distance R does not vary with Γ , which means that the maximum entropy distribution is not closer to the target distribution than any other distribution compatible with the observable values. Our results are confirmed by Monte Carlo sampling of the version space for small system sizes ($N \leq 10$).

Keywords Probabilistic inference · Maximum entropy principle · Replica method

1 Introduction

Multi-components and strongly interacting systems, in physics and beyond, may show complex behaviours eluding simple quantitative modeling. A common strategy to describe such systems is to define probability distributions over the space of their configurations. The task is at first sight daunting. The number of unknown probabilities scales as the dimension of the configuration space, and is enormous, generally exponentially large in the number of the degrees of freedom defining the system. The selection of one probability distribution among the multitude of possibilities may be done in a Bayesian way. A class of parametrized model distributions is considered, and a good choice of the parameters is sought, *e.g.* which maximizes the likelihood of the observed data. Statisticians are generally interested in understanding how learning proceeds, that is, in the speed of convergence to the target distribution (assumed to be in the parametrized class, and to be consistent with one value of the parameters) as more and more data are made available. Yet, though the number of parameters defining the class of distributions may

T. Obuchi

Department of Computational Intelligence and Systems Science, Tokyo Institute of Technology, Yokohama, Japan
E-mail: obuchi@sp.dis.titech.ac.jp

S. Cocco

Laboratoire de Physique Statistique de l’Ecole Normale Supérieure, affilié au CNRS & à l’Université Pierre et Marie Curie, Paris, France

R. Monasson

Laboratoire de Physique Théorique de l’Ecole Normale Supérieure, affilié au CNRS & à l’Université Pierre et Marie Curie, Paris, France

be very large and the inference problem may be computationally very hard, the classes of parametrized distributions are generally extremely small compared to the plethora of possible distributions over the configuration space. While some classes of distributions may appear more adequate to represent the data or more amenable to computations, those criteria are largely arbitrary.

The Maximum Entropy (ME) principle is an alternative approach, rejecting the notion of arbitrariness. The ME principle may be informally stated as follows: among all possible distributions compatible with what is known of the data, choose the one with ME. It was proposed as an alternative foundation of statistical mechanics [1,2]; as an illustration, the Boltzmann distribution in the canonical ensemble may be found back as the distribution with ME under the sole knowledge of the average value of the energy. The ME principle is supported by information theory [3], which argues that any other distribution would be too constrained, and, in fact, reflect additional properties of the data [4]. An alternative and somewhat colourful formulation of this argument was given in [5], where the ME distribution was shown to emerge as the most frequent distribution which an uninformed operator (monkeys in [5]) would find, upon repeated and unbiased trials compatible with the observations on the data. In other words, given an unbiased prior over the set of all possible distributions over the configuration space, the ME distribution is the most likely one compatible with our knowledge of the data. Furthermore the ME distribution enjoys important and valuable properties, *e.g.* a weak sensitivity to measurement errors [6]. To the practitioner, the most compelling argument in support of the ME distribution may actually come from the successful applications to experimental data, see for instance [7] for a clear presentation regarding biological data.

The purpose of the present work is to compare the performances of the ME distribution with the ones of the other distributions compatible with the data. While the literature on the ME principle has a rich history, we are not aware of works attempting to carry out this comparison in a quantitative and rather general setting. We consider a set of N variables, taking binary values ± 1 (the generalization to a larger number of states would be straightforward, as long as it remains finite). Any distribution over the space of configurations is entirely characterized by 2^N probabilities, which are non-negative real numbers summing to unity. Each distribution may therefore be seen as a point in the 2^N -dimensional simplex. We pick up one point in this simplex, hereafter referred to as the target distribution. Next we consider a set of M observables, each of which may be any polynomial function of the variables, and compute their average values over the target distribution. The distributions in the simplex compatible with those data, *i.e.* such that the observables have the same average within some prescribed accuracy, are called admissible. The set of admissible distributions, called version space, contains the target distribution, the ME distributions, and many others, such as the center-of-mass distribution, which is the flat average over all admissible distributions. Our objective is to compute the main geometrical features of the version space, *e.g.* its volume, the distances from the target to the ME or to the center-of-mass distributions, the average inter-distribution distance,

To give a precise meaning to those quantities in a mathematically tractable framework we will assume that observables are drawn from a simple statistical ensemble. More precisely the values taken by the observables are assumed to be random and uncorrelated across the 2^N configurations. This assumption is not meant to be realistic. In most real applications, indeed, observables reflect the low-order statistics of the configurations, *e.g.* the value of the first variable, the product of the fifth and seventh variables, Such observables vary very smoothly over the configuration space, as they depend only a small (compared to N) number of the configuration variables, and may be adequate to provide information about smooth distributions. Our hypothesis can be considered as worst-case-like in the following sense. The inference of the target distribution from the average values of the observables may be recast as the problem of reconstructing a 2^N -dimensional non-negative vector from the knowledge of its scalar products with M vectors (corresponding to the observables). When those vectors are randomly chosen the scalar products are typically very small, of the order of $2^{-N/2}$, and are weakly informative about the direction of the target vector. In this pessimistic setting we expect that the number of scalar products (observables) necessary to reconstruct the target vector with accuracy will be of the order of the number of relevant components of the target vector, that is, of the order of the exponential of the entropy of the target distribution. While this statement is correct we will see that some important features of the target distribution are correctly inferred even with a much smaller number M of observables. In particular, the probabilities of configurations with large values of the target distribution are learned with a limited number of data, a phenomenon connected to the onset of phase transitions in the learning process.

An important aspect of our framework is that it allows us to bias the measure over the version space, in order to boost the distributions with large entropies. The magnitude of this entropic bias may be

continuously tuned from zero (uniform measure over the version space) to infinity, which amounts to selecting the ME distribution alone. We study how the distance between the target distribution and center-of-mass distribution varies with the bias (for a fixed number M of observables). Our main result is that this distance does not depend on the bias, showing that the ME entropy is not better than any other distribution randomly picked up in the version space. While this result is valid for any target distribution in the case of random observables we do not expect it to apply to the more realistic case where both the observables and the target distribution are smooth.

All our results are derived within the replica symmetric framework when N and M go to infinity at fixed ratio $\alpha = (\log M)/N$, and are therefore non rigorous. We have, however, checked the local stability of our replica-symmetric solution against replica-symmetry-breaking fluctuations, the so-called replicon modes. Our results are therefore self-consistent and we expect them to be correct for large N, M . In addition we have designed a Monte Carlo algorithm to sample the version space, and applied it to small system sizes ($N \leq 10$). Remarkably, simulations show only weak finite-size effects compared to our large- N calculations; a good qualitative (and sometimes even quantitative) agreement with our analytical predictions is found.

The paper is organized to be accessible to the reader not interested in the details of our calculation. In sec. 2, we present the necessary definitions and notations. An overview of our results, free of technicalities, is given in sec. 3. All technical details and calculations are reported in sec. 4 and in the Appendix. We present the sampling algorithm and the results of our numerical simulations in sec. 5. Conclusions can be found in sec. 6.

2 Definitions and notations

2.1 Target distribution

Let us consider a system consisting of N Ising spins, with configurations $\mathbf{s} = \{s_i = \pm 1\}_{i=1}^N$. The probability distributions of the system configurations, hereafter called target distribution, is denoted by $\hat{p}_{\mathbf{s}}$. We consider large-size systems, $N \gg 1$, and write

$$\hat{p}_{\mathbf{s}} \doteq e^{-N\omega_{\mathbf{s}}}, \quad (1)$$

where the rate $\omega_{\mathbf{s}} > 0$ of the configuration probability $\hat{p}_{\mathbf{s}}$ is introduced, and the symbol \doteq stands for equality in the leading exponential-in- N term. It is convenient to introduce the entropy σ of the rates $\omega_{\mathbf{s}}$,

$$\sigma(\omega) = \lim_{\epsilon \rightarrow 0^+} \lim_{N \rightarrow \infty} \frac{1}{N} \log \left[\text{nb. of configurations } \mathbf{s} \text{ such that } e^{-N\omega} \leq \hat{p}_{\mathbf{s}} < e^{-N(\omega-\epsilon)} \right], \quad (2)$$

The entropy curve $\sigma(\omega)$ has some remarkable features in standard physical systems. First, it is convex and bounded from above and below. Secondly, the maximum of the curve is $\sigma_0 = \log 2$; the corresponding value of ω is denoted by ω_0 . Thirdly, the curve lies below the $\sigma = \omega$ line, and is tangent to this line at $\omega_1 (\leq \omega_0)$. The value $\sigma_1 = \sigma(\omega_1)$ is the entropy per spin of the target distribution

$$\sigma_1 = \lim_{N \rightarrow \infty} -\frac{1}{N} \sum_{\mathbf{s}} \hat{p}_{\mathbf{s}} \log \hat{p}_{\mathbf{s}}. \quad (3)$$

More generally, given a real number k , we define ω_k , σ_k , and ℓ_k through

$$\left. \frac{d\sigma}{d\omega} \right|_{\omega=\omega_k} = k, \quad \sigma_k \equiv \sigma(\omega_k), \quad \ell_k \equiv k\omega_k - \sigma_k. \quad (4)$$

Note that the range of values of k such that ω_k is well-defined is generally bounded. These quantities characterize the dominant contribution to the k th moment of \hat{p} :

$$\sum_{\mathbf{s}} (\hat{p}_{\mathbf{s}})^k = \int d\omega e^{N(\sigma(\omega) - k\omega)} \doteq e^{N(\sigma_k - k\omega_k)} = e^{-N\ell_k}. \quad (5)$$

As an illustration of the properties above we show in Fig. 1 the entropy curve of the independent spin model (ISM) defined as

$$\hat{p}_s^{\text{ISM}} = \frac{1}{(2 \cosh H)^N} \exp \left(H \sum_{i=1}^N s_i \right), \quad (6)$$

for $H = 0.5$. The value of H will not change throughout this paper.




Fig. 1 Entropy σ as a function of ω for the independent spin model in eq. (6) with $H = 0.5$. The entropy curve is obtained through a parametric representation $\omega(m) = \log(2 \cosh H) - Hm$, $\sigma(m) = \log 2 - \frac{1}{2}(1+m) \log(1+m) - \frac{1}{2}(1-m) \log(1-m)$, where m is the magnetization per spin, ranging from -1 to $+1$. It is easy to show that $\omega_k = \omega(\tanh(kH))$.

2.2 Observables, version space, and maximum entropy distribution

Let \mathbf{v} be an observable, taking value v_s when the system configuration is \mathbf{s} . Observables and probability distributions can be seen as vectors in a 2^N -dimensional space, with components labelled by the configurations \mathbf{s} . We assume that measurements give us access to the average value of the observable over the target distribution,

$$\sum_{\mathbf{s}} v_{\mathbf{s}} \hat{p}_{\mathbf{s}} \equiv \mathbf{v} \cdot \hat{\mathbf{p}}, \quad (7)$$

which may be simply written as the scalar product of the vectors attached to the observable and to the target distribution. Suppose we have made M measurements corresponding to M observables \mathbf{v}^μ ($\mu = 1, \dots, M$). An ‘admissible’ probability vector, \mathbf{p} , compatible with all the measurements is such that

$$(\mathbf{p} - \hat{\mathbf{p}}) \cdot \mathbf{v}^\mu = 0, \quad \forall \mu = 1, \dots, M. \quad (8)$$

We hereafter use the term ‘constraints’ to refer to eq. (8) or to the attached observables \mathbf{v}^μ . The set of vectors satisfying eq. (8), together with the normalization and the non-negativity conditions

$$\sum_{\mathbf{s}} p_{\mathbf{s}} = 1 \quad \text{and} \quad p_{\mathbf{s}} \geq 0, \forall \mathbf{s}, \quad (9)$$

defines the version space.

Each distribution \mathbf{p} in the version space is characterized by its Shannon entropy,

$$S(\mathbf{p}) = - \sum_{\mathbf{s}} p_{\mathbf{s}} \log p_{\mathbf{s}}, \quad (10)$$

which may be interpreted as an estimate of the logarithm of the effective number of configurations under the distribution. The maximum entropy (ME) distribution, \mathbf{p}_{ME} , is the distribution maximizing eq. (10) in the version space, *i.e.* under the constraints listed in eq. (8) and in eq. (9). Using Lagrange multipliers η^μ to enforce those constraints the ME distribution may be formally written as

$$p_{\mathbf{s}}^{\text{ME}} = C \exp \left(\sum_{\mu} \eta^{\mu} v_{\mathbf{s}}^{\mu} \right), \quad (11)$$

where C is a normalization constant. As an illustration, if we consider the set of the N single-spin observables, $v_{\mathbf{s}} = s_i$, $1 \leq i \leq N$, and of the $N(N-1)/2$ two-spin observables $v_{\mathbf{s}} = s_i s_j$, $1 \leq i < j \leq N$, we recover the well-known result that the Ising model is the ME model given the average values of those observables. The Lagrange multipliers η coincide with, respectively, the N fields and the $N(N-1)/2$ pairwise couplings acting on the spins.

2.3 Measures over the space of distributions

In realistic situations involving certain measurement noise, it may be beneficial not to perfectly reproduce the given average value by avoiding overfitting. To take into account such a flexibility, we introduce a Gaussian-like measure with variance E over the vector space of \mathbf{p} :

$$\rho(\mathbf{p} | \{\mathbf{v}^{\mu}\}_{\mu=1}^M, \hat{\mathbf{p}}) = \frac{1}{V} \prod_{\mathbf{s}} \theta(p_{\mathbf{s}}) \delta \left(\sum_{\mathbf{s}} p_{\mathbf{s}} - 1 \right) \exp \left\{ -\frac{1}{2E} \sum_{\mu=1}^M (\mathbf{v}^{\mu} \cdot (\mathbf{p} - \hat{\mathbf{p}}))^2 \right\}, \quad (12)$$

where $\theta(x)$ is the step function, equal to 1 if $x \geq 0$ and to 0 if $x < 0$, and $\delta(x)$ is the Dirac delta function. We will call E tolerance hereafter. The denominator V is defined to make sure that the measure ρ is normalized:

$$V(E, \{\mathbf{v}^{\mu}\}_{\mu=1}^M, \hat{\mathbf{p}}) = \int_0^\infty \prod_{\mathbf{s}} dp_{\mathbf{s}} \delta \left(\sum_{\mathbf{s}} p_{\mathbf{s}} - 1 \right) \exp \left\{ -\frac{1}{2E} \sum_{\mu=1}^M (\mathbf{v}^{\mu} \cdot (\mathbf{p} - \hat{\mathbf{p}}))^2 \right\}. \quad (13)$$

This normalization factor measures the volume of ‘admissible’ probability vectors given the constraints eq. (8). In this probabilistic setting we will loosely use the term ‘version space’ to refer to the set of distributions \mathbf{p} associated to ‘large’ measure values $\rho(\mathbf{p})$. Note that, while $\rho(\mathbf{p})$ defines the joint-measure of the 2^N -configuration probabilities, we will also consider below the marginal measure for a single-configuration probability,

$$\rho_{\mathbf{s}}(p_{\mathbf{s}} | \{\mathbf{v}^{\mu}\}_{\mu=1}^M, \hat{\mathbf{p}}) = \int_0^\infty \prod_{\mathbf{t}(\neq \mathbf{s})} dp_{\mathbf{t}} \rho(\mathbf{p} | \{\mathbf{v}^{\mu}\}_{\mu=1}^M, \hat{\mathbf{p}}). \quad (14)$$

It is possible to consider other measures of the space of \mathbf{p} for the purpose of studying the performances of the ME distribution. To favor the probability vectors \mathbf{p} with large Shannon entropies $S(\mathbf{p})$, see eq. (10), it is natural to introduce the new measure

$$\rho(\mathbf{p} | \Gamma, \{\mathbf{v}^{\mu}\}_{\mu=1}^M, \hat{\mathbf{p}}) \propto \rho(\mathbf{p} | \{\mathbf{v}^{\mu}\}_{\mu=1}^M, \hat{\mathbf{p}}) e^{\Gamma S(\mathbf{p})}, \quad (15)$$

where Γ is the strength of this entropic bias. The normalization of this new measure ρ implicitly defines the following expression for the volume in the presence of an entropic bias,

$$V(\Gamma, E, \{\mathbf{v}^{\mu}\}_{\mu=1}^M, \hat{\mathbf{p}}) = \int_0^\infty \prod_{\mathbf{s}} dp_{\mathbf{s}} \delta \left(\sum_{\mathbf{s}} p_{\mathbf{s}} - 1 \right) \exp \left\{ -\frac{1}{2E} \sum_{\mu=1}^M (\mathbf{v}^{\mu} \cdot (\mathbf{p} - \hat{\mathbf{p}}))^2 + \Gamma S(\mathbf{p}) \right\}. \quad (16)$$

For $\Gamma = 0$ we recover the ‘unbiased’ measure in eq. (12), while the limit $\Gamma \rightarrow \infty$ retains the ME distribution only.

2.4 Randomization of observables

Hereafter we assume that the M constraints $\{\mathbf{v}^\mu\}_{\mu=1}^M$ are randomly and independently chosen from the following Gaussian distribution over the space of 2^N -dimensional vectors \mathbf{v} :

$$P(\mathbf{v}) = \prod_s \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}v_s^2}. \quad (17)$$

As mentioned in the introduction we do not pretend that this assumption is realistic. Indeed in practice, observables are often chosen to be low-order moments of the target distribution such as magnetizations or pairwise spin correlations, which is quite distinct from the Gaussian distribution above.

In contrast to the constraints $\{\mathbf{v}^\mu\}_\mu$, which are randomly drawn, we do not choose any particular statistical ensemble for the target distribution $\hat{\mathbf{p}}$; The only properties of $\hat{\mathbf{p}}$ we need for the analysis is the existence of the entropy curve discussed above. We stress that these constraints $\{\mathbf{v}^\mu\}_\mu$, appearing in the measure ρ over the distribution space, are quenched random variables, drawn once for all.

Since each constraint contributing a multiplicative factor $e^{-(\mathbf{v}^\mu \cdot (\mathbf{p} - \hat{\mathbf{p}}))^2 / 2E}$ to the volume V in eq. (16), we expect that the logarithm of this volume will be self-averaging in the large- M limit. We will therefore calculate the average value of $\log V$ over the constraints, using the replica method.

3 Overview of results

We hereafter report the main outcomes of our replica calculation in an informal way. All technical details are postponed to sec. 4.

3.1 Order parameters: interpretation and scaling with N

Given the set of constraints $\{\mathbf{v}^\mu\}_{\mu=1}^M$ and the target distribution $\hat{\mathbf{p}}$ we may draw a schematic picture of the version space such as the one shown in Fig. 2. In addition to the target distribution, two distributions of interest in the version space are the ME distribution, \mathbf{p}^{ME} , and the center-of-mass distribution, $\langle \mathbf{p} \rangle$, where the angular brackets denote the average over the measure ρ :

$$\langle \mathbf{p} \rangle = \int d\mathbf{p} \, \rho(\mathbf{p}) \, \mathbf{p}. \quad (18)$$

The following quantities are useful to characterize the ‘distances’ between the distributions in the version space, see Fig. 2:

$$R(\{\mathbf{v}^\mu\}_\mu, \hat{\mathbf{p}}) \equiv \sum_s (\langle p_s \rangle - \hat{p}_s)^2, \quad D(\{\mathbf{v}^\mu\}_\mu, \hat{\mathbf{p}}) \equiv \sum_s (\langle p_s^2 \rangle - \langle p_s \rangle^2). \quad (19)$$

R measures the L_2 -squared distance between the target and the center-of-mass distributions. D gives the size of the squared fluctuations of \mathbf{p} around the center of mass. We will also consider hereafter

$$Q(\{\mathbf{v}^\mu\}_\mu, \hat{\mathbf{p}}) \equiv R + D = \sum_s \langle (p_s - \hat{p}_s)^2 \rangle, \quad (20)$$

which measures the averaged square distance between \mathbf{p} and $\hat{\mathbf{p}}$ (Fig. 2).

D, Q, R are intimately related to the statistics of the generalization error, that is, the error on the prediction of the average value of a new observable. Assume we have a measure ρ over the version space defined from a set of constraints $\{\mathbf{v}^\mu\}_{\mu=1}^M$. Let us now consider a new observable, \mathbf{v}' . The error on the average value of this observable \mathbf{v}' computed with a distribution \mathbf{p} is simply given by

$$\Delta(\mathbf{p}) = \mathbf{v}' \cdot (\mathbf{p} - \hat{\mathbf{p}}). \quad (21)$$

Suppose now that \mathbf{p} is randomly chosen according to measure ρ and that the components v'_s are independent and normal random variables, with zero means and unit variances. Let us denote the average over \mathbf{v}' by $[\dots]_{\mathbf{v}'}$. It is easy to show that

$$Q = [\langle \Delta^2 \rangle]_{\mathbf{v}'}, \quad R = [\langle \Delta \rangle^2]_{\mathbf{v}'} = [\langle \Delta \rangle^2]_{\mathbf{v}'} - [\langle \Delta \rangle]_{\mathbf{v}'}^2, \quad D = [\langle \Delta^2 \rangle - \langle \Delta \rangle^2]_{\mathbf{v}'}. \quad (22)$$




Fig. 2 Schematic view of the space of distributions. The large circle represents the version space of all admissible probability vectors, see eq. (8), which depends on the constraints $\{\mathbf{v}^\mu\}_{\mu=1}^M$. The shaded area represents the typical fluctuations of \mathbf{p} around the center-of-mass distribution $\langle \mathbf{p} \rangle$, of magnitude \sqrt{D} . The distance between the target distribution $\hat{\mathbf{p}}$ and $\langle \mathbf{p} \rangle$ is \sqrt{R} , while \sqrt{Q} is the square root of the averaged squared distance between \mathbf{p} and $\hat{\mathbf{p}}$. These three order parameters measure the lengths of the sides of the rectangular triangle shown in the figure, see eq. (20). The ME distribution, \mathbf{p}^{ME} , lies inside the version space.

Hence, Q represents the averaged square error on a new observable, R quantifies the observable-to-observable squared fluctuations of the error, and D measures the distribution-to-distribution squared fluctuations of the error. The squared error on the prediction of the average value of the new observable, Q , is the sum of two contributions. The first one, R , is due to the choice of observable. The second one, D , reflects the distribution-to-distribution fluctuations with the version space measure.

To understand the scaling of D, Q, R with N let us consider the simple case of no constraint at all, $M = 0$, and no bias on the entropy, $\Gamma = 0$. In this case, due to the permutation symmetry over the 2^N configurations, all the configuration probabilities p_s obey the same marginal distribution, $\rho_s(p_s)$, see eq. (14). It is easy to convince oneself, that ρ_s becomes a pure exponential when $N \gg 1$, with average value $\langle p_s \rangle = 2^{-N}$. Indeed, the volume in eq. (13) is given by

$$V = \int_0^\infty \prod_s dp_s \delta\left(\sum_s p_s - 1\right) = \int_{-\infty}^\infty d\Lambda e^\Lambda \int_0^\infty \prod_s (dp_s e^{-\Lambda p_s}) = \int_{-\infty}^\infty d\Lambda \frac{e^\Lambda}{\Lambda^{2^N}}, \quad (23)$$

where we have used an integral representation of the Dirac delta function. This can be directly integrated but we here use the saddle-point method, valid for large N for later convenience. The saddle point is located at

$$\Lambda = 2^N, \quad (24)$$

giving $\log V \doteq 2^N(1 - N \log 2)$, which agrees with the leading behaviour of the volume of the 2^N -dimensional simplex, $V = 1/[(2^N)!]$. In addition we see from eq. (14) and eq. (23) that ρ_s is an exponential distribution with mean value $1/\Lambda = 2^{-N}$, as announced above. Hence,

$$D \doteq e^{Nd}, \quad Q \doteq e^{Nq}, \quad R \doteq e^{Nr}, \quad \Lambda \doteq e^{N\lambda}, \quad (25)$$

with $q = r = -\ell_2$, see eq. (4), and $d = -\lambda = \log 2$. Our calculation with the replica method shows that, in the presence of constraints, *i.e.* for $M \geq 1$, the exponential-in- N scaling of D, Q, R, Λ will still hold after averaging over the constraints, even if a bias Γ over the entropy is imposed. Note that the rates d, q, r, λ will then depend on M and Γ .

3.2 Description of the learning process

We first focus on the effect of increasing the number of measured observables, M , and set the entropic bias Γ to zero. The effect on non-zero biases will be reported in sec. 3.3. We assume in sec. 3.2.1 and sec. 3.2.2 that the tolerance E is negligible; the dependence of the results on the value of E is exposed in sec. 3.2.3.

3.2.1 The learning edge

One major result of the replica calculation is that the marginal measure over the single-configuration probabilities, $\rho_{\mathbf{s}}$, eq. (14), may have two distinct behaviours, sketched in Fig. 3:

- The marginal measure over the probability $p_{\mathbf{s}}$ of a configuration \mathbf{s} having a *large* value $\hat{p}_{\mathbf{s}}$ in the target distribution is concentrated at a value close to this target probability:

$$\rho_{\mathbf{s}}(p_{\mathbf{s}}) \approx A e^{-B(p_{\mathbf{s}} - \hat{p}_{\mathbf{s}} - \delta_{\mathbf{s}})^2/2}, \quad (26)$$

where the values of A and B depend on the parameters M, E , and will be specified later. The corresponding curve of $\rho_{\mathbf{s}}$ is shown in Fig. 3, left. The shift $\delta_{\mathbf{s}}$ is small compared to the target probability $\hat{p}_{\mathbf{s}}$, as we will see in sec. 4. In other words, the configuration probability $\hat{p}_{\mathbf{s}}$ is correctly ‘learned’ from the values of the constraints.

- The marginal measure over the probability $p_{\mathbf{s}}$ of a configuration \mathbf{s} having a *small* value $\hat{p}_{\mathbf{s}}$ in the target distribution is a decaying exponential:

$$\rho_{\mathbf{s}}(p_{\mathbf{s}}) \approx \frac{1}{A'} e^{-p_{\mathbf{s}}/A'}. \quad (27)$$

The corresponding curve of $\rho_{\mathbf{s}}$ is shown in Fig. 3, right. The average value A' of the configuration probability $p_{\mathbf{s}}$ does not depend on $\hat{p}_{\mathbf{s}}$ in the dominant scaling with N . A' is a function of the parameters M, E only, and will be specified later. In other words, the configuration probability is not at all inferred.




Fig. 3 Schematic pictures of the marginal measure $\rho_{\mathbf{s}}(p_{\mathbf{s}})$ for large (left) and small (right) configuration probabilities $\hat{p}_{\mathbf{s}}$.

The concepts of *large* and *small* target probabilities are defined as follows:

$$\hat{p}_{\mathbf{s}} \doteq e^{-N\omega_{\mathbf{s}}} \text{ is large if } \omega_{\mathbf{s}} < \hat{\omega}, \text{ and is small if } \omega_{\mathbf{s}} > \hat{\omega}. \quad (28)$$

The boundary $\hat{\omega}$ is hereafter referred to as the *learning edge*, as it separates the set of probabilities $p_{\mathbf{s}}$ into the ones which can be correctly learned and the ones which remain essentially unknown, despite the knowledge of the observable values. Its value depends on the parameters M and E . The representative curve of the learning edge $\hat{\omega}$ is an increasing function of the rate

$$\alpha = \frac{\log M}{N}, \quad (29)$$

and is shown in Fig. 4 for the ISM and a negligible tolerance E . We find qualitative changes of the curve $\hat{\omega}(\alpha)$, taking place at critical values of the ratio α and corresponding to the onset of phase transitions in the learning process.




Fig. 4 Learning edge $\hat{\omega}$ as a function of α for the ISM (6) with $H = 0.5$ and for small E .

3.2.2 Phase transitions and typical distances between distributions

Given the value of the learning edge $\hat{\omega}$ we define, according to eq. (28), the sets of configurations \mathbf{s} having large and small probabilities as, respectively, L and S . Our replica calculation shows that the order parameters Q and R are given by, to the dominant order in N ,

$$Q \doteq R \doteq \sum_{\mathbf{s} \in S} \hat{p}_{\mathbf{s}}^2. \quad (30)$$

If M is small, so is the learning edge $\hat{\omega}$, most configurations are found in S . This implies $\sum_{\mathbf{s} \in S} \hat{p}_{\mathbf{s}}^2 \doteq \sum_{\mathbf{s}} \hat{p}_{\mathbf{s}}^2 = e^{-N\ell_2}$. As M increases, the learning edge grows, and reaches $\hat{\omega} = \omega_2$. For larger values of M , the dominant point $\omega = \omega_2$ is now not in $\mathbf{s} \in S$ and hence the rates q and r change from $-\ell_2$. This leads non-analyticities in Q and R , and the transition point is at $\hat{\omega} = \omega_2$. Similarly, a switch of the dominant term in the normalization condition $1 = \sum_{\mathbf{s} \in S} \hat{p}_{\mathbf{s}} + \sum_{\mathbf{s} \in L} \hat{p}_{\mathbf{s}}$ from the S to the L occurs at $\hat{\omega} = \omega_1$. The Lagrange multiplier Λ is non analytic at this point, which defines a second phase transition.

Hence, we have three distinct phases, separated by specific values of the learning edge. We label the phases by ‘I’ when $\hat{\omega} < \omega_2$, ‘II’ when $\omega_2 \leq \hat{\omega} < \omega_1$, and ‘III’ when $\omega_1 \leq \hat{\omega}$. The value of the ratio α corresponding to the transition point $\hat{\omega} = \omega_2$ is denoted as α_2 , and the one corresponding to $\hat{\omega} = \omega_1$ is called α_1 . In phases I and II with small E , $\hat{\omega}$ turns out to be equal to α . Thus $\alpha_2 = \omega_2$ and $\alpha_1 = \omega_1$.

To get insights about the typical distances between the distributions of interest in the version space we plot the order parameters q, r and d of the ISM in Fig. 5. Those order parameters continuously change at the transition points. At $\alpha = \log 2$, the learning edge becomes equal to ω_0 ; q and r reach the same value as d . The agreement between r and d , the distance of the center of mass to the target distribution and the typical fluctuations, implies that the volume largely shrinks and scales with N in a different manner at this point. Our calculation, restricted to the leading order in the volume, is not informative for larger ratios, *i.e.* $\alpha > \log 2$.

3.2.3 Effect of the tolerance E

Let us now consider the case of a non-negligible tolerance E . We give the corresponding phase diagram in Fig. 6. It is natural to scale the tolerance as $E \doteq e^{N\epsilon}$. As ϵ grows from very negative values, we expect that the quality of the inference becomes worse, as the version space include more distributions, which do not exactly reproduce the target values of the observables. Indeed, for large values of the tolerance E , there emerge new phases where the learning edge decreases as ϵ grows. We call these phases with large




Fig. 5 Order parameters q, r, d as functions of the learning edge $\hat{\omega}$ for the ISM (6) with $H = 0.5$ and for small E . The critical values of ω are $\omega_0 \approx 0.81$, $\omega_1 \approx 0.58$, and $\omega_2 \approx 0.43$, and the corresponding entropies are $\sigma_0 = \log 2 \approx 0.69$, $\sigma_1 \approx 0.58$, and $\sigma_2 \approx 0.37$, respectively.




Fig. 6 Phase diagram in the α - ϵ plane in the absence of any entropic bias.

tolerance I_{LT} , II_{LT} , and III_{LT} , in agreement with the denomination chosen above based on the different ranges of possible values for the learning edge.

The transition from small to large tolerances takes place at $E \doteq D$. This result is easy to interpret: deviations in the values of observables from the target ones smaller than the scale of the intrinsic fluctuations D will be masked by those fluctuations, and cannot degrade the inference performances. The results shown in sec. 3.2.1 and sec. 3.2.2 are therefore valid as long as $\epsilon < d$.

In the present study, we do not consider any measurement error: the measured values of the observables correspond exactly to the projection of the observable vectors onto the target distribution. If measurement errors were considered, the role of the tolerance parameter could possibly change; an appropriate amount of tolerance E , of the order of the measurement error, would lead to better inference by avoiding overfitting.

3.3 Effects of the entropic bias

We now consider the effect of an entropic bias $\Gamma > 0$ on the learning performances. As Γ grows, the measure ρ gives more and more weight to the distributions \mathbf{p} with large entropies. In the $\Gamma \rightarrow \infty$ limit ρ

singles out the ME distribution. The Γ -dependence of the quantities of interest, in particular, the learning edge, is of key importance. Similarly to the other parameters we assume that the entropic bias scales $\Gamma = e^{N\gamma}$.

We first consider the case of negligible $E \approx 0$. Our replica calculation shows that the order parameters Q, R and the learning edge $\hat{\omega}$ do not depend at all on the value $\Gamma \geq 0$. An immediate consequence is that the ME distribution is not closer to the target distribution than any randomly picked up distribution (with the unbiased measure) in the version space. Nevertheless the presence of an entropic bias affects the fluctuations of \mathbf{p} , measured by the order parameter D . For $\alpha < \alpha_1$, our replica calculation shows that $D \doteq 2^{-N}$ for small Γ , while $D \doteq \Gamma^{-1}$ for large Γ . The transition between these two scalings takes place at $\gamma_c = \log 2$. The same transition for D is observed when $\alpha > \alpha_1$, though the transition point γ_c now depends on α . The order parameter Λ shows non-analyticities at those transition points, contrary to Q, R , and the learning edge, which remain unaffected as mentioned above. The non-analyticities in D and Λ discriminate the phases for $\gamma \leq \gamma_c$ from the ones for $\gamma > \gamma_c$; we denote the phases with large entropic bias by I_{ME}, II_{ME} , and III_{ME} , in agreement with the denomination based on the ranges of values of the learning edge. The corresponding phase diagram is shown in Fig. 7, left. On the right panel of the same figure, the order parameter d is plotted against γ for the ISM with $H = 0.5$ and $\alpha < \alpha_1$ as an illustration.




Fig. 7 Left: Phase diagram in the α - γ plane for very small tolerance E . The fluctuation order parameter $D \doteq e^{Nd}$ depends on γ in the phases with large entropic biases, I_{ME}, II_{ME} , and III_{ME} , while it keeps the same value as for $\Gamma = 0$ in I, II, and III. Right: order parameter d as a function of γ for the ISM and $\alpha < \alpha_1$. A non-analyticity appears at $\gamma = \gamma_c (= \log 2)$, which signals the onset of the phase with large entropic bias for $\gamma > \gamma_c$.

The phases above change when the value of the tolerance $E \doteq e^{N\epsilon}$ is not negligible any longer. As in the $\Gamma = 0$ case the tolerance matters if $E > D$. As the fluctuations scale as $D \doteq \Gamma^{-1}$ for large Γ with $\alpha < \alpha_1$, this implies that a regime takes place for $\gamma > -\epsilon$. This large tolerance, large entropic bias regime shows unusual properties. For instance the learning edge becomes smaller and smaller, that is, the learning performances become worse and worse, as Γ grows, in contrast to the case of negligible $E \approx 0$ where the learning edge has no dependence on Γ . This seemingly-strange behaviour is actually expected. We see from eq. (16) that, for very large Γ and finite fixed E , the constraints become irrelevant. Hence, in this limit, the measure concentrates around the uniform distribution $p_s^{ME} = 1/2^N$, maximizing the entropy irrespective of the constraints. To get an meaningful ME distribution, the tolerance E must be enough small compared to the fluctuations D governed by Γ in the large Γ limit. In conclusion, this large tolerance, large entropic bias regime is irrelevant and will not be studied further.

4 Analytical calculations

In this section, we present our replica calculation of the average value of the logarithm of the volume V over the constraints. We introduce the notations

$$\mathrm{Tr}_{\mathbf{p}}(\cdots) \equiv \int_0^\infty \prod_{\mathbf{s}} dp_{\mathbf{s}} \delta\left(\sum_{\mathbf{s}} p_{\mathbf{s}} - 1\right) (\cdots), \quad (31)$$

and $Dz \equiv dz e^{-\frac{1}{2}z^2}/\sqrt{2\pi}$. We also write undefined integrals as a shorthand notation for the domains of integration $[-\infty, \infty]$ or $[-i\infty, i\infty]$.

4.1 Replica calculation of the average logarithm of the volume

We start by unraveling the squared terms in the exponent of volume (16) through auxiliary Gaussian integrations, or the so-called Hubbard-Stratonovich transformations:

$$V = \int \prod_{\mu=1}^M Dz_{\mu} \mathrm{Tr}_{\mathbf{p}} \exp \left\{ i \sum_{\mu} \sum_{\mathbf{s}} \frac{z_{\mu}}{\sqrt{E}} v_{\mathbf{s}}^{\mu} (p_{\mathbf{s}} - \hat{p}_{\mathbf{s}}) + \Gamma S(\mathbf{p}) \right\}. \quad (32)$$

According to the replica method we calculate the n th moment of the volume of $n \in \mathbb{N}$

$$\begin{aligned} [V^n] &= \prod_{a=1}^n \left(\int \prod_{\mu} Dz_{\mu}^a \mathrm{Tr}_{\mathbf{p}^a} \right) \left[\exp \left\{ i \sum_{\mu} \sum_{\mathbf{s}} \sum_a \frac{z_{\mu}^a}{\sqrt{E}} v_{\mathbf{s}}^{\mu} (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}}) + \Gamma \sum_a S(\mathbf{p}^a) \right\} \right] \\ &= \prod_{a=1}^n \left(\int \prod_{\mu} Dz_{\mu}^a \mathrm{Tr}_{\mathbf{p}^a} \right) \exp \left\{ - \sum_{\mu} \sum_{a,b} \frac{z_{\mu}^a z_{\mu}^b}{2E} Q_{ab} + \Gamma \sum_a S(\mathbf{p}^a) \right\} \\ &= \prod_{a=1}^n \mathrm{Tr}_{\mathbf{p}^a} \exp \left\{ - \frac{M}{2} \log \det \left(1 + \frac{Q}{E} \right) + \Gamma \sum_a S(\mathbf{p}^a) \right\}, \end{aligned} \quad (33)$$

where the square brackets denote the average over the observables $\{\mathbf{v}^{\mu}\}_{\mu=1}^M$ as in sec. 3.1, and we define

$$Q_{ab} = \sum_{\mathbf{s}} (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}})(p_{\mathbf{s}}^b - \hat{p}_{\mathbf{s}}). \quad (34)$$

To perform the integrations $\mathrm{Tr}_{\mathbf{p}^a}$, we make use of the following identity

$$\begin{aligned} 1 &= \int \prod_{a \leq b} dQ_{ab} \prod_{a \leq b} \delta \left(Q_{ab} - \sum_{\mathbf{s}} (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}})(p_{\mathbf{s}}^b - \hat{p}_{\mathbf{s}}) \right) \\ &= C \int \prod_{a \leq b} dQ_{ab} dQ'_{ab} \exp \left\{ \frac{1}{2} \sum_{a \leq b} Q_{ab} Q'_{ab} - \frac{1}{2} \sum_{a \leq b} \sum_{\mathbf{s}} Q'_{ab} (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}})(p_{\mathbf{s}}^b - \hat{p}_{\mathbf{s}}) \right\}, \end{aligned} \quad (35)$$

where the normalization constant C is irrelevant and will be discarded hereafter. Similarly, we rewrite $\mathrm{Tr}_{\mathbf{p}^a}$ as

$$\mathrm{Tr}_{\mathbf{p}^a} = \int d\Lambda_a \int_0^\infty \prod_{\mathbf{s}} dp_{\mathbf{s}}^a e^{-\Lambda_a \sum_{\mathbf{s}} (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}})}, \quad (36)$$

where we used the normalization identity $1 = \sum_{\mathbf{s}} \hat{p}_{\mathbf{s}}$. We then obtain

$$\begin{aligned} [V^n] &= \int \prod_{a \leq b} dQ_{ab} dQ'_{ab} \int \prod_a d\Lambda_a \int_0^\infty \prod_a \prod_{\mathbf{s}} dp_{\mathbf{s}}^a \\ &\times \exp \left(- \sum_a \Lambda_a \sum_{\mathbf{s}} (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}}) - \frac{1}{2} \sum_{a \leq b} \sum_{\mathbf{s}} Q'_{ab} (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}})(p_{\mathbf{s}}^b - \hat{p}_{\mathbf{s}}) + \Gamma \sum_a S(\mathbf{p}^a) \right) \\ &\times \exp \left\{ \frac{1}{2} \sum_{a \leq b} Q_{ab} Q'_{ab} - \frac{M}{2} \log \det \left(1 + \frac{Q}{E} \right) \right\}. \end{aligned} \quad (37)$$

The logarithm of $[V^n]$ can be approximated by the saddle-point value of

$$\phi(n) \equiv \frac{1}{2} \sum_{a \leq b} Q_{ab} Q'_{ab} - \frac{M}{2} \log \det \left(1 + \frac{Q}{E} \right) + \sum_{\mathbf{s}} \log \Theta_{\mathbf{s}}, \quad (38)$$

over the matrices Q and Q' , and where

$$\Theta_{\mathbf{s}} = \int_0^\infty \prod_a dp_{\mathbf{s}}^a \exp \left(- \sum_a \Lambda_a (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}}) - \frac{1}{2} \sum_{a \leq b} Q'_{ab} (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}})(p_{\mathbf{s}}^b - \hat{p}_{\mathbf{s}}) - \Gamma \sum_a p_{\mathbf{s}}^a \log p_{\mathbf{s}}^a \right). \quad (39)$$

We now assume that the order parameter matrices are invariant under permutation symmetry of the replica indices, and write

$$Q_{ab} = R + (Q - R)\delta_{ab}, \quad Q'_{ab} = -2R' + (Q' + R')\delta_{ab}, \quad \Lambda_a = \Lambda, \quad (40)$$

where δ_{ab} is the Kronecker delta. Thus

$$\frac{1}{2} \sum_{a \leq b} Q_{ab} Q'_{ab} = \frac{1}{2} n Q (Q' - R') - \frac{1}{2} n(n-1) R R', \quad (41)$$

$$\log \det \left(1 + \frac{Q}{E} \right) = \log \left(1 + \frac{Q + (n-1)R}{E} \right) + (n-1) \log \left(1 + \frac{Q-R}{E} \right), \quad (42)$$

and,

$$\frac{1}{2} \sum_{a \leq b} Q'_{ab} (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}})(p_{\mathbf{s}}^b - \hat{p}_{\mathbf{s}}) = \frac{1}{2} Q' \sum_a (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}})^2 - \frac{1}{2} R' \left(\sum_a (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}}) \right)^2. \quad (43)$$

Using the Hubbard-Stratonovich transformation again eq. (39) becomes

$$\Theta_{\mathbf{s}} = \int Dz \int_0^\infty \prod_a dp_{\mathbf{s}}^a e^{-(\Lambda - z\sqrt{R'}) (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}}) - \frac{1}{2} Q' (p_{\mathbf{s}}^a - \hat{p}_{\mathbf{s}})^2 - \Gamma p_{\mathbf{s}}^a \log p_{\mathbf{s}}^a} \equiv \int Dz X_{\mathbf{s}}^n. \quad (44)$$

Hence,

$$\begin{aligned} \phi(n) &= \frac{1}{2} n Q (Q' - R') - \frac{1}{2} n(n-1) R R' \\ &\quad - \frac{M}{2} \left\{ \log \left(1 + \frac{Q + (n-1)R}{E} \right) + (n-1) \log \left(1 + \frac{Q-R}{E} \right) \right\} + \sum_{\mathbf{s}} \log \int Dz X_{\mathbf{s}}^n. \end{aligned} \quad (45)$$

We finally obtain the expression for $[\log V] = \lim_{n \rightarrow 0} \phi(n)/n$,

$$[\log V] = \frac{1}{2} Q (Q' - R') + \frac{1}{2} R R' - \frac{M}{2} \left\{ \frac{R}{E + Q - R} + \log \left(1 + \frac{Q-R}{E} \right) \right\} + \sum_{\mathbf{s}} \int Dz \log X_{\mathbf{s}}, \quad (46)$$

where

$$X_{\mathbf{s}} = \int_0^\infty dp e^{-(\Lambda - z\sqrt{R'}) (p - \hat{p}_{\mathbf{s}}) - \frac{1}{2} Q' (p - \hat{p}_{\mathbf{s}})^2 - \Gamma p \log p}. \quad (47)$$

The integration over p in eq. (47) has to be done with care. We assume that the conjugated order parameters, Q' and R' , obey an exponential scaling with N ,

$$Q' \doteq e^{Nq'}, \quad R' \doteq e^{Nr'}. \quad (48)$$

Therefore, all the parameters appearing in the integration diverge or vanish when $N \rightarrow \infty$. We need to find out which order parameters diverge or not, consistently with the equations of state derived below. These procedures require involved case analyses. The outcome is that the equations of state are consistent, to the leading order in N , if the following conditions are met:

$$0 < \lambda < q' \leq 2\lambda, \quad 0 < \lambda < r' \leq 2\lambda, \quad q = r < 0, \quad \epsilon < -\gamma \leq 0. \quad (49)$$

4.2 $\Gamma = 0$ case

In this case, the integral in eq. (47) can be evaluated in two different ways. The first one is more direct, but the latter one is useful to clarify the physical significance of the solutions and to treat the finite- Γ case.

4.2.1 Direct integration and expansion

The first strategy to evaluate X_s is to directly integrate eq. (47) in $\Gamma = 0$. The result is

$$X_s = \int_0^\infty dp e^{-(\Lambda - z\sqrt{R'})(p - \hat{p}_s) - \frac{1}{2}Q'(p - \hat{p}_s)^2} = \sqrt{\frac{2\pi}{Q'}} e^{\frac{1}{2}\frac{(\Lambda - z\sqrt{R'})^2}{Q'}} H\left(y_s - z\sqrt{\frac{R'}{Q'}}\right), \quad (50)$$

where we put $y_s = (\Lambda - \hat{p}_s Q')/\sqrt{Q'}$ and we define the complementary error function

$$H(y) = \int_y^\infty Dz. \quad (51)$$

Thus

$$\begin{aligned} [\log V] &= \frac{1}{2}Q(Q' - R') + \frac{1}{2}RR' - \frac{M}{2} \left\{ \frac{R}{E + Q - R} + \log\left(1 + \frac{Q - R}{E}\right) \right\} \\ &+ 2^N \left(\frac{\Lambda^2 + R'}{2Q'} + \frac{1}{2}\log 2\pi - \frac{1}{2}\log Q' \right) + \sum_s \int Dz \log H\left(y_s - z\sqrt{\frac{R'}{Q'}}\right). \end{aligned} \quad (52)$$

The definition of y_s allows us to introduce the learning edge in a natural way. Under assumption (49), $y_s = (\Lambda - \hat{p}_s Q')/\sqrt{Q'}$ diverges. If the target probability $\hat{p}_s \doteq e^{-N\omega_s}$ is small enough, *i.e.* if $\omega_s > q' - \lambda$, the dominant term in y_s is $\Lambda/\sqrt{Q'}$, which goes to $+\infty$ in the thermodynamic limit. On the contrary we find $y_s \rightarrow -\infty$ if $\omega_s < q' - \lambda$. This sharp difference defines the learning edge and the corresponding entropy as

$$\hat{\omega} = q' - \lambda, \quad \hat{\sigma} = \sigma(\hat{\omega}) \quad (53)$$

Furthermore, based on assumption (49), $|y_s|$ increases faster with $N \gg 1$ than $\sqrt{R'/Q'}$, entailing that the argument of the complementary error function H in eq. (50) is dominated by y_s . Thus, the complementary error function can be replaced with its asymptotic behavior, depending on the sign of y_s ,

$$H(y) \rightarrow \begin{cases} \frac{1}{\sqrt{2\pi}} \frac{1}{y} e^{-\frac{1}{2}y^2}, & (y \rightarrow \infty) \\ 1, & (y \rightarrow -\infty) \end{cases}. \quad (54)$$

Using this, we get

$$\begin{aligned} &\sum_s \int Dz \log H\left(y_s - z\sqrt{\frac{R'}{Q'}}\right) \\ &= \int Dz \sum_{s \in S} \left\{ -\frac{1}{2} \left(y_s - z\sqrt{\frac{R'}{Q'}} \right)^2 - \frac{1}{2} \log 2\pi - \log \left(y_s - z\sqrt{\frac{R'}{Q'}} \right) \right\} + \sum_{s \in L} 0. \end{aligned} \quad (55)$$

The logarithmic term in the above formula can be expanded according to eq. (49) as

$$\log \left(y_s - z\sqrt{\frac{R'}{Q'}} \right) \approx \log \Lambda - \frac{1}{2} \log Q' - \frac{\hat{p}_s Q' + \sqrt{R'} z}{\Lambda} - \frac{1}{2} \frac{(\hat{p}_s Q' + \sqrt{R'} z)^2}{\Lambda^2}. \quad (56)$$

Linear terms with respect to z vanish upon Gaussian integration. We are left with

$$\begin{aligned} [\log V] &\approx \frac{1}{2}Q(Q' - R') + \frac{1}{2}RR' - \frac{M}{2} \left\{ \frac{R}{E + Q - R} + \log \left(1 + \frac{Q - R}{E} \right) \right\} \\ &+ e^{N\hat{\sigma}} \left(\frac{\Lambda^2 + R'}{2Q'} + \frac{1}{2} \log 2\pi - \frac{1}{2} \log Q' \right) - 2^N \left(\log \Lambda - \frac{1}{2} \frac{R'}{\Lambda^2} \right) \\ &+ \left(\Lambda + \frac{Q'}{\Lambda} \right) \sum_{\mathbf{s} \in S} \hat{p}_{\mathbf{s}} - \frac{1}{2} \left(Q' - \frac{Q'^2}{\Lambda^2} \right) \sum_{\mathbf{s} \in S} \hat{p}_{\mathbf{s}}^2, \end{aligned} \quad (57)$$

where we have used

$$2^N - \sum_{\mathbf{s} \in S} 1 = \sum_{\mathbf{s} \in L} 1 \doteq e^{N\hat{\sigma}}, \quad \sum_{\mathbf{s} \in S} 1 \doteq 2^N. \quad (58)$$

These formulas are correct as long as $\hat{\omega} \leq \omega_0$.

4.2.2 Saddle-point approximation

Our second method to evaluate $X_{\mathbf{s}}$ in eq. (50) is based on the saddle-point approximation. Writing the exponent as $f_{\mathbf{s}}(p) = -\left(\Lambda - z\sqrt{R'}\right)(p - \hat{p}_{\mathbf{s}}) - \frac{1}{2}Q'(p - \hat{p}_{\mathbf{s}})^2$ in eq. (50) and differentiating with respect to p we get the solution of the saddle-point equation $f'_{\mathbf{s}}(p^*) = 0$:

$$p^* = \hat{p}_{\mathbf{s}} - \frac{\Lambda - z\sqrt{R'}}{Q'}. \quad (59)$$

As we know from the previous calculation $\frac{1}{N} \log(\Lambda/Q')$ determines the learning edge. If $\hat{p}_{\mathbf{s}} > \Lambda/Q'$ the dominant contribution to p^* is $\hat{p}_{\mathbf{s}}$ itself and becomes positive, implying the saddle point is in the feasibility domain:

$$X_{\mathbf{s}} \approx \int_{-\infty}^{\infty} dx \, e^{f_{\mathbf{s}}(p^*) - \frac{1}{2}|f''_{\mathbf{s}}(p^*)|x^2} = e^{f_{\mathbf{s}}(p^*)} \sqrt{\frac{2\pi}{|f''_{\mathbf{s}}(p^*)|}} = e^{\frac{1}{2} \frac{(\Lambda - z\sqrt{R'})^2}{Q'}} \sqrt{\frac{2\pi}{Q'}}. \quad (60)$$

This result coincides with eq. (50) with $H(y) \rightarrow 1$, $y \rightarrow -\infty$. To check the validity of the saddle-point approximation we can compare the peak location $p^* \doteq \hat{p}_{\mathbf{s}}$ to the width of the Gaussian, $\sigma = 1/\sqrt{|f''_{\mathbf{s}}(p^*)|} = 1/\sqrt{Q'}$. Under assumption (49) we see that $\omega_{\mathbf{s}} \leq \hat{\omega} = q' - \lambda < q'/2$ and thus $p^* > \sigma$ to dominant exponential-in- N order. The peak height $f_{\mathbf{s}}(p^*)$ also diverges rapidly with N , and the saddle-point approximation is justified.

In the case of configurations $\mathbf{s} \in S$ with small probabilities, $\omega_{\mathbf{s}} > \hat{\omega}$, the dominant term to the right hand side of eq. (59) is $-\Lambda/Q'$; p^* becomes negative and lies outside the domain of integration. Hence, the true saddle-point value is $p^* = 0$. We expand $f_{\mathbf{s}}(p)$ around $p = 0$ up to the first order, and approximate the integral as

$$X_{\mathbf{s}} \approx \int_0^{\infty} dp \, e^{f_{\mathbf{s}}(0) + f'_{\mathbf{s}}(0)p} = \frac{e^{(\Lambda - z\sqrt{R'})\hat{p}_{\mathbf{s}} - \frac{1}{2}Q'\hat{p}_{\mathbf{s}}^2}}{\Lambda - Q'\hat{p}_{\mathbf{s}} - z\sqrt{R'}}. \quad (61)$$

This is again identical to eq. (50) in the limit $H(y) \approx e^{-y^2/2}/(\sqrt{2\pi}y)$ with $y = y_{\mathbf{s}} - z\sqrt{R'}/Q'$.

In addition, the saddle-point calculations above give us the expressions for the marginal measure $\rho_{\mathbf{s}}$ given in sec. 3.2. Indeed, the marginal measure of the probability of the configuration \mathbf{s} reads

$$\rho_{\mathbf{s}}(p_{\mathbf{s}}) = \frac{1}{X_{\mathbf{s}}} e^{f_{\mathbf{s}}(p_{\mathbf{s}})}. \quad (62)$$

Therefore $\rho_{\mathbf{s}}$ is approximately an exponentially-decaying function (27) for small target probabilities, and a Gaussian centered close to the corresponding target probability (26) for large target probabilities.

4.2.3 Equations of state and their solutions

Here we derive the equations of state (EOSs). Taking the derivatives of eq. (57) with respect to the order parameters, we get

$$R' = \frac{MR}{(E + Q - R)^2}, \quad Q' = \frac{M}{E + Q - R}, \quad (63)$$

$$Q = \left(1 - 2\frac{Q'}{\Lambda^2}\right) \sum_{s \in S} \hat{p}_s^2 - \frac{2}{\Lambda} \sum_{t \in S} \hat{p}_s + e^{N\hat{\sigma}} \left(\frac{\Lambda^2 + R' + Q'}{Q'^2}\right), \quad (64)$$

$$D = Q - R = \frac{2^N}{\Lambda^2} + \frac{e^{N\hat{\sigma}}}{Q'}, \quad (65)$$

$$e^{N\hat{\sigma}} \frac{\Lambda}{Q'} - \frac{2^N}{\Lambda} + \left(1 - \frac{Q'}{\Lambda^2}\right) \sum_{s \in S} \hat{p}_s - \frac{Q'^2}{\Lambda^3} \sum_{s \in S} \hat{p}_s^2 = 0. \quad (66)$$

The order parameter D defined in sec. 3 naturally appears. According to assumption (49), we know that all the terms with the factor Q'/Λ^2 are not dominant in eqs. (63-66); the same statement applies to the terms with the factor $e^{N\hat{\sigma}}$ in eqs. (64-66). Matching the dominant contributions to the remaining terms we obtain

$$Q \doteq \sum_{s \in S} \hat{p}_s^2, \quad \frac{2^N}{\Lambda} \doteq \sum_{s \in S} \hat{p}_s, \quad (67)$$

$$D \doteq \frac{2^N}{\Lambda^2}. \quad (68)$$

The physical meaning of the first equation is clear. The order parameter $Q = \sum_s \langle (p_s - \hat{p}_s)^2 \rangle$ quantifies the average squared distance of a distribution \mathbf{p} (chosen with measure ρ) to the target distribution. The contributions coming from the large probabilities are negligible because, as we have seen in sec. 4.2.2, the marginal measures for large probabilities are centered close to the target values, which makes $\sum_{s \in L} \langle (p_s - \hat{p}_s)^2 \rangle$ very small. However, for the configurations for small target probabilities, $\langle p_s^2 \rangle$ is much smaller than \hat{p}_s^2 . Thus Q is, to the leading order in N , equal to $\sum_{s \in S} \hat{p}_s^2$. We see, in addition, from eq. (67)

that Λ ensures the normalization condition, and is equal to 2^N as long as most configurations are in the set S . According to eq. (67), we find

$$q = r = \begin{cases} -\ell_2 & (\hat{\omega} < \omega_2) \\ \hat{\sigma} - 2\hat{\omega} & (\omega_2 \leq \hat{\omega}) \end{cases}, \quad \lambda = \begin{cases} \log 2 & (\hat{\omega} < \omega_1) \\ \log 2 - \hat{\sigma} + \hat{\omega} & (\omega_1 \leq \hat{\omega}) \end{cases}, \quad d = \log 2 - 2\lambda. \quad (69)$$

The other order parameters q', r' and $\hat{\omega}$ are computed accordingly.

Small- E case. We can ignore E when it is smaller than $D = Q - R$. In this case we get the EOSs for the parameters q' and r' :

$$r' = \alpha + r - 2d, \quad q' = \alpha - d, \quad (70)$$

and the learning edge is self-consistently determined by eq. (53). Solving theses equations, we obtain

Phase I ($\hat{\omega} < \omega_2$):

$$\lambda = \log 2, \quad \hat{\omega} = \alpha, \quad q' = \alpha + \log 2, \quad q = r = -\ell_2, \quad r' = \alpha + 2\log 2 - \ell_2, \quad d = -\log 2. \quad (71)$$

Phase II ($\omega_2 \leq \hat{\omega} < \omega_1$):

$$\lambda = \log 2, \quad \hat{\omega} = \alpha, \quad q' = \alpha + \log 2, \quad q = r = \sigma(\alpha) - 2\alpha, \\ r' = \sigma(\alpha) - \alpha + 2\log 2, \quad d = -\log 2. \quad (72)$$

Phase III ($\omega_1 < \hat{\omega}$):

$$\begin{aligned}\lambda &= \log 2 + \hat{\omega} - \alpha, \quad \hat{\omega} = \sigma^{-1}(\alpha), \quad q' = \log 2 + 2\hat{\omega} - \alpha, \quad q = r = \sigma(\hat{\omega}) - 2\hat{\omega}, \\ r' &= 2(\log 2 + \hat{\omega} - \alpha) = 2\lambda, \quad d = -\log 2 + 2\hat{\sigma} - 2\hat{\omega}.\end{aligned}\tag{73}$$

The critical ratio separating phases I and II is given by $\alpha_2 = \omega_2$, while the critical ratio associated to the transition from II to III is $\alpha_1 = \omega_1$. We can easily check that those solutions satisfy the assumptions (49), and are therefore self-consistent.

Large- E case. If $E \doteq e^{N\epsilon}$ is larger than $D \doteq e^{Nd}$, i.e. if $\epsilon > d$, the EOSs of the parameters q' and r' are modified as

$$r' = \alpha + r - 2\epsilon, \quad q' = \alpha - \epsilon.\tag{74}$$

These large-tolerance EOSs become valid beyond the critical values $\hat{\epsilon}(\alpha)$:

$$\hat{\epsilon}(\alpha) = -\log 2 \text{ (Phase I, II)}, \quad \hat{\epsilon}(\alpha) = -\log 2 + 2\hat{\sigma}(\alpha) - 2\hat{\omega}(\alpha) \text{ (Phase III)}.\tag{75}$$

To distinguish the solutions above from the ones of the small- E case we denote phase I with large tolerance by I_{LT} as explained in sec. 3.2.3. Phases II_{LT} and III_{LT} are defined in the same way. The solutions of the large-tolerance case become

Phase I_{LT} ($\hat{\omega} < \omega_2$):

$$\begin{aligned}\lambda &= \log 2, \quad \hat{\omega} = \alpha - \epsilon - \log 2, \quad q' = \alpha - \epsilon = \hat{\omega} + \log 2, \quad q = r = -\ell_2, \\ r' &= \alpha - 2\epsilon - \ell_2 = 2\log 2 + 2\hat{\omega} - \ell_2 - \alpha, \quad d = -\log 2.\end{aligned}\tag{76}$$

Phase II_{LT} ($\omega_2 \leq \hat{\omega} < \omega_1$):

$$\begin{aligned}\lambda &= \log 2, \quad \hat{\omega} = \alpha - \epsilon - \log 2, \quad q' = \alpha - \epsilon = \hat{\omega} + \log 2, \quad q = r = \hat{\sigma} - 2\hat{\omega}, \\ r' &= \alpha + \hat{\sigma} - 2\hat{\omega} - 2\epsilon = \hat{\sigma} + 2\log 2 - \alpha, \quad d = -\log 2.\end{aligned}\tag{77}$$

Phase III_{LT} ($\omega_1 \leq \hat{\omega}$):

$$\begin{aligned}\lambda &= \log 2 + \hat{\omega} - \hat{\sigma}, \quad \hat{\sigma} - 2\hat{\omega} = \log 2 + \epsilon - \alpha, \quad q' = \alpha - \epsilon, \quad q = r = \hat{\sigma} - 2\hat{\omega}, \\ r' &= \alpha + \hat{\sigma} - 2\hat{\omega} - 2\epsilon = 2\log 2 - \alpha - \hat{\sigma} + 2\hat{\omega}, \quad d = -\log 2 + 2\hat{\sigma} - 2\hat{\omega}.\end{aligned}\tag{78}$$

It is instructive to examine the validity of the condition $\epsilon > d$ in phase III_{LT} . To do so we need to know the rate of increase of $\hat{\sigma} - \hat{\omega}$ with ϵ at fixed α . Differentiating the equation determining the learning edge in eqs. (78), we get

$$\frac{\partial(\hat{\sigma} - \hat{\omega})}{\partial\epsilon} = 1 + \frac{\partial\hat{\omega}}{\partial\epsilon}.\tag{79}$$

This equation can be written as

$$-1 \leq \frac{\partial\hat{\omega}}{\partial\epsilon} = \left(\frac{\partial\hat{\sigma}}{\partial\hat{\omega}} - 2 \right)^{-1} \leq -\frac{1}{2}.\tag{80}$$

The inequalities come from the condition $\partial\hat{\sigma}/\partial\hat{\omega} < 1$, which holds in phase III. This implies that d increases more slowly than ϵ itself in III_{LT} , since $\partial d/\partial\epsilon = 2(1 + \partial\hat{\omega}/\partial\epsilon) \leq 1$. Thus the necessary condition $\epsilon > d$ is satisfied in phase III_{LT} , as it should. This condition will be useful for the study of the stability of the RS ansatz.

4.3 Large- Γ case

4.3.1 Saddle-point approximation

We use the saddle-point approximation in the integral defining X_s (47) as in sec. 4.2.2. We call $f_s(p)$ the exponent in the integral, and write down the derivatives

$$f_s(p) = -\frac{1}{2}Q'(p - \hat{p}_s)^2 - (\Lambda - z\sqrt{R'})(p - \hat{p}_s) - \Gamma p \log p, \quad (81)$$

$$f'_s(p) = -Q'(p - \hat{p}_s) - (\Lambda - z\sqrt{R'}) - \Gamma(\log p + 1), \quad (82)$$

$$f''_s(p) = -Q' - \frac{\Gamma}{p}. \quad (83)$$

The main difference with the $\Gamma = 0$ case is the presence of the term $p \log p$, singular at $p = 0$. As a result the functional behavior around $p = 0$ changes completely, and there is always a peak in $[0, \infty]$. This implies that the saddle-point approximation can be applied, irrespective of the target probability value. The saddle-point equation $f_s(p^*) = 0$ can be written as

$$p^* = \hat{p}_s - \frac{\Lambda - z\sqrt{R'} + \Gamma(1 + \log p^*)}{Q'}. \quad (84)$$

Let us make the following assumption, correct when Γ is large enough,

$$\frac{\Lambda}{\Gamma} = O(N). \quad (85)$$

This relation, in turns, defines the meaning of ‘large’ Γ . This assumption, combined with eq. (49), gives us the value of the learning edge from eq. (84): $\hat{\omega} = q' - \lambda = q' - \gamma$. For configurations with large target probabilities the dominant term is \hat{p}_s in eq. (84), and an iterative substitution yields

$$p_s^* \approx \hat{p}_s - \frac{\Lambda - z\sqrt{R'} + \Gamma(1 + \log \hat{p}_s)}{Q'}, \quad (s \in L). \quad (86)$$

Expanding $f_s(p)$ up to the second order, we may evaluate X_s .

For configurations corresponding to small target probabilities, the saddle point is very different. We *a priori* postulate the expression of the solution

$$p_s^* \approx \tilde{p} + \Delta_s, \quad (s \in S), \quad (87)$$

where the dominant term \tilde{p} is

$$\tilde{p} = e^{-1 - (\Lambda - Q'\tilde{p}_s - z\sqrt{R'})/\Gamma} < \hat{p}_s, \quad (\forall s \in L). \quad (88)$$

Sec. A gives a reasoning of this form. \tilde{p} is required to take this value to satisfy the normalization condition. According to the saddle-point calculation above we have

$$1 = \left\langle \sum_s p_s \right\rangle \approx \sum_{s \in S} \tilde{p} + \sum_{s \in L} \hat{p}_s. \quad (89)$$

For small learning edge values, $\hat{\omega} < \omega_1$, configurations with small probabilities dominate, implying that $\sum_{s \in S} \tilde{p} \doteq 2^N \tilde{p} \approx 1 \Rightarrow \tilde{p} \doteq 2^{-N}$. Thus, \tilde{p} is automatically determined by the normalization condition. Moreover, the correction term Δ_s in eq. (87) is determined by expanding eq. (84) with respect to Δ_s to the first order and solving the resulting equation. We find

$$\Delta_s = -\frac{Q'\tilde{p}^2}{Q'\tilde{p} + \Gamma}. \quad (90)$$

This expansion is valid if $\tilde{p} > |\Delta_s|$, yielding an trivial inequality $1 > Q'\tilde{p}/(Q'\tilde{p} + \Gamma)$.

4.3.2 Equations of state and validity of the saddle-point approximation

It is convenient to compute the derivatives with respect to the order parameters of the expression in eq. (46). The result is

$$Q = \sum_{\mathbf{s}} \int Dz \left\langle (p - \hat{p}_{\mathbf{s}})^2 \right\rangle_{X_{\mathbf{s}}}, \quad (91)$$

$$D = Q - R = \sum_{\mathbf{s}} \int Dz \left\{ \left\langle (p - \hat{p}_{\mathbf{s}})^2 \right\rangle_{X_{\mathbf{s}}} - \langle p - \hat{p}_{\mathbf{s}} \rangle_{X_{\mathbf{s}}}^2 \right\}, \quad (92)$$

$$0 = \sum_{\mathbf{s}} \int Dz \langle p - \hat{p}_{\mathbf{s}} \rangle_{X_{\mathbf{s}}} \quad (93)$$

where we define

$$\langle (\dots) \rangle_{X_{\mathbf{s}}} = \frac{1}{X_{\mathbf{s}}} \int_0^\infty dp (\dots) e^{f_{\mathbf{s}}(p)}. \quad (94)$$

The EOSs for R' and Q' are identical to eq. (63), and are thus omitted above. The average value of p is replaced with the saddle-point value $p_{\mathbf{s}}^*$. We need to quantify the fluctuations around the saddle-point to estimate the terms in eq. (92), that is,

$$\left\langle (p - \hat{p}_{\mathbf{s}})^2 \right\rangle_{X_{\mathbf{s}}} - \langle p - \hat{p}_{\mathbf{s}} \rangle_{X_{\mathbf{s}}}^2 \approx \frac{e^{f_{\mathbf{s}}(p_{\mathbf{s}}^*)} \int dx x^2 e^{-\frac{f_{\mathbf{s}}''(p_{\mathbf{s}}^*)}{2} x^2}}{e^{f_{\mathbf{s}}(p_{\mathbf{s}}^*)} \int dx e^{-\frac{|f_{\mathbf{s}}''(p_{\mathbf{s}}^*)|}{2} x^2}} = \frac{1}{|f_{\mathbf{s}}''(p_{\mathbf{s}}^*)|} \doteq \begin{cases} Q'^{-1} & (\mathbf{s} \in L) \\ \frac{\tilde{p}}{T} & (\mathbf{s} \in S) \end{cases}. \quad (95)$$

Note that, for $\mathbf{s} \in S$, we can write $p_{\mathbf{s}}^* - \hat{p}_{\mathbf{s}} \doteq \tilde{p} - \hat{p}_{\mathbf{s}}$. These considerations lead us to

$$\begin{aligned} Q &= \sum_{\mathbf{s} \in S} \int Dz (\tilde{p} - \hat{p}_{\mathbf{s}})^2 + \sum_{\mathbf{s} \in L} \int Dz \left\{ \left(-\frac{\Lambda + \Gamma + \Gamma \log \hat{p}_{\mathbf{s}}}{Q'} - z \frac{\sqrt{R'}}{Q'} \right)^2 \right\} \\ &\doteq 2^N \tilde{p}^2 - 2\tilde{p} \sum_{\mathbf{s} \in S} \hat{p}_{\mathbf{s}} + \sum_{\mathbf{s} \in S} \hat{p}_{\mathbf{s}}^2 + e^{N\hat{\sigma}} \left(\frac{\Lambda + \Gamma}{Q'} \right)^2 \left\{ \left(1 - \frac{N\hat{\omega}\Gamma}{\Lambda + \Gamma} \right)^2 + \frac{R'}{(\Lambda + \Gamma)^2} \right\}. \end{aligned} \quad (96)$$

$$D = \sum_{\mathbf{s} \in S} \int Dz \frac{\tilde{p}}{\Gamma} + \sum_{\mathbf{s} \in L} \int Dz (Q')^{-1} \doteq 2^N \frac{\tilde{p}}{\Gamma} + \frac{e^{N\hat{\sigma}}}{Q'}, \quad (97)$$

$$\begin{aligned} 0 &= \sum_{\mathbf{s} \in S} \int Dz (\tilde{p} - \hat{p}_{\mathbf{s}}) + \sum_{\mathbf{s} \in L} \int Dz \left(-\frac{\Lambda + \Gamma + \Gamma \log \hat{p}_{\mathbf{s}}}{Q'} - z \frac{\sqrt{R'}}{Q'} \right) \\ &\doteq 2^N \tilde{p} - \sum_{\mathbf{s} \in S} \hat{p}_{\mathbf{s}} - e^{N\hat{\sigma}} \left(\frac{\Lambda + \Gamma}{Q'} \right) \left(1 - \frac{N\hat{\omega}\Gamma}{\Lambda + \Gamma} \right). \end{aligned} \quad (98)$$

The dominant terms turn out to be

$$Q \doteq \sum_{\mathbf{s} \in S} \hat{p}_{\mathbf{s}}^2, \quad (99)$$

$$D \doteq 2^N \frac{\tilde{p}}{\Gamma}, \quad (100)$$

$$2^N \tilde{p} \doteq \sum_{\mathbf{s} \in S} \hat{p}_{\mathbf{s}}, \quad (101)$$

with the corresponding exponents,

$$q = r = \begin{cases} -\ell_2 & (\hat{\omega} < \omega_2) \\ \hat{\sigma} - 2\hat{\omega} & (\omega_2 \leq \hat{\omega}) \end{cases}, \quad d = \begin{cases} -\gamma & (\hat{\omega} < \omega_1) \\ \hat{\sigma} - \hat{\omega} - \gamma & (\omega_1 \leq \hat{\omega}) \end{cases}, \quad (102)$$

$$\lambda = \gamma + \frac{\log N}{N} + \frac{1}{N} \begin{cases} \log \log 2 & (\hat{\omega} < \omega_1) \\ \log(\log 2 + \hat{\omega} - \hat{\sigma}) & (\omega_1 \leq \hat{\omega}) \end{cases}. \quad (103)$$

Small- E case. Again we can neglect $E(\ll D)$. The EOSs for q' and r' and the equation for the learning edge are fully identical to eq. (70) and eq. (53). Each phase is then characterized as follows:

Phase I_{ME} ($\hat{\omega} < \omega_2$):

$$\lambda = \gamma + \frac{\log(N \log 2)}{N}, \quad \hat{\omega} = \alpha, \quad q' = \alpha + \gamma, \quad q = r = -\ell_2, \quad r' = \alpha + 2\gamma - \ell_2, \quad d = -\gamma. \quad (104)$$

Phase II_{ME} ($\omega_2 \leq \hat{\omega} < \omega_1$):

$$\begin{aligned} \lambda &= \gamma + \frac{\log(N \log 2)}{N}, \quad \hat{\omega} = \alpha, \quad q' = \alpha + \gamma, \quad q = r = \sigma(\alpha) - 2\alpha, \\ r' &= \sigma(\alpha) - \alpha + 2\gamma, \quad d = -\gamma. \end{aligned} \quad (105)$$

Phase III_{ME} ($\omega_1 < \hat{\omega}$):

$$\begin{aligned} \lambda &= \gamma + \frac{\log(N(\log -\alpha + \hat{\omega})2)}{N}, \quad \hat{\omega} = \sigma^{-1}(\alpha), \quad q' = \hat{\omega} + \gamma, \quad q = r = \alpha - 2\hat{\omega}, \\ r' &= 2\gamma, \quad d = \hat{\sigma} - \hat{\omega} - \gamma. \end{aligned} \quad (106)$$

It is easy to check that the assumption eq. (49) is satisfied by these solutions. Note that the learning edge does not depend on γ .

We now examine the validity of the saddle-point approximation for $X_{\mathbf{s}}$, to determine when the solutions above are correct. To do so we compare the width of the Gaussian with the location of its peak value, and check that the peak height diverges. For example, in the phase I_{ME},

Peak location:

$$p_{\mathbf{s}}^* \doteq \begin{cases} \hat{p}_{\mathbf{s}} \doteq e^{-N\omega_{\mathbf{s}}}, & (\mathbf{s} \in L) \\ \tilde{p} \doteq e^{-N \log 2}, & (\mathbf{s} \in S) \end{cases}. \quad (107)$$

Height:

$$f_{\mathbf{s}}(p_{\mathbf{s}}^*) \doteq \begin{cases} e^{N(\gamma + \hat{\omega} - 2\omega_{\mathbf{s}})}, & (\mathbf{s} \in L) \\ e^{N(\gamma - \log 2)}, & (\mathbf{s} \in S) \end{cases}. \quad (108)$$

Width:

$$\sigma_{\mathbf{s}} = |f_{\mathbf{s}}''(p_{\mathbf{s}}^*)|^{-1/2} = \sqrt{\frac{p_{\mathbf{s}}^*}{Q'p_{\mathbf{s}}^* + \Gamma}} \doteq \begin{cases} (Q')^{-1/2} \doteq e^{-N\frac{\gamma + \hat{\omega}}{2}}, & (\mathbf{s} \in L) \\ \sqrt{\frac{\tilde{p}}{\Gamma}} \doteq e^{-N\frac{\gamma + \log 2}{2}}, & (\mathbf{s} \in S) \end{cases}. \quad (109)$$

According to these relations the saddle-point approximation is correct if $\gamma \geq \log 2 \equiv \gamma_c$. The same approach can be applied to II_{ME} and III_{ME}, with the resulting critical values: $\gamma_c = \log 2$ in II_{ME} and $\gamma_c = \log 2 + \hat{\omega} - \hat{\sigma}$ for III_{ME}. These critical values are consistent with eq. (85), as Λ is a continuous function of Γ .

This continuous change also implies that, in the $0 < \Gamma < \Gamma_c = e^{N\gamma_c}$ region, all solutions coincide with the ones found for $\Gamma = 0$. Consider for instance the peak value of the marginal measure over $p_{\mathbf{s}}$ for $\mathbf{s} \in S$. Assuming the order parameters are given by their expressions in the $\Gamma = 0$ case, we see that the condition eq. (49) is satisfied, which implies that eq. (87) is valid. The saddle point for $\mathbf{s} \in S$, $\tilde{p} = e^{-1 - (\Lambda - Q'\tilde{p} - z\sqrt{R'})/\Gamma}$, decays in a double-exponential manner with respect to N , as Λ is exponentially larger than Γ . Similarly, the peak height $f_{\mathbf{s}}(p_{\mathbf{s}}^*)$ rapidly goes to zero for $\mathbf{s} \in S$. The peak location, \tilde{p} , decreases faster than the width $\sigma_{\mathbf{s}} \doteq \sqrt{\tilde{p}/\Gamma}$. As a consequence the functional form $e^{f_{\mathbf{s}}(p)}$ rapidly converges to a pure exponential distribution, as in the $\Gamma = 0$ case.

Large- E case. The EOSs for q' and r' coincide with eq. (74). The critical values of E are determined by $E \doteq D$:

$$\hat{\epsilon}(\alpha, \gamma) = -\gamma \text{ (Phase I, II)}, \quad \hat{\epsilon}(\alpha, \gamma) = \hat{\sigma}(\alpha) - \hat{\omega}(\alpha) - \gamma \text{ (Phase III)}. \quad (110)$$

New solutions can appear for $\epsilon > \hat{\epsilon}(\alpha, \gamma)$, but should be discarded as they violate the condition $E < \Gamma$, see discussion in sec. 3.3.

4.4 Stability of the replica symmetry

We study the de Almeida-Thouless stability of the replica-symmetric (RS) Ansatz, see eq. (40), against fluctuations in the replica space [8]. Detailed calculations are reported in the Appendix, sec. B. The outcome of the calculation is the following stability condition against replicon fluctuations:

$$\frac{M}{(E + Q - R)^2} \left\{ \sum_s \int Dz \left(\frac{\partial^2 \log X_s}{\partial \Lambda^2} \right)^2 \right\} \leq 1. \quad (111)$$

The term in the brackets is asymptotically equivalent to $e^{N\hat{\sigma}}/Q'^2$. This can be shown in the small- Γ case ($\Gamma \leq \Gamma_c$) using eqs. (60,61), and in the large- Γ case ($\Gamma > \Gamma_c$) with the relation $\partial^2 \log X_s / \partial \Lambda^2 = \langle (p - \hat{p})^2 \rangle_{X_s} - \langle (p - \hat{p}) \rangle_{X_s}^2$ and eq. (95). From eq. (63), we obtain a transparent interpretation of the stability condition

$$\frac{M}{(E + Q - R)^2} \left(\frac{e^{N\hat{\sigma}}}{Q'^2} \right) = \frac{e^{N\hat{\sigma}}}{M} \doteq e^{N(\hat{\sigma} - \alpha)} \leq 1. \quad (112)$$

Therefore the RS solution becomes unstable if the number of large-probability configurations exceeds the number of constraints. Inserting the solutions for $\hat{\sigma}$ in the small E regime, see eqs. (71-73) and eqs. (104-106), we find that, irrespective of the value of Γ , the RS ansatz is stable in phases I and II but is only marginally stable in phase III. Marginal stability means here that $\hat{\sigma}$ is equal to α in phase III. Our calculation, limited to the leading order in N , cannot decide whether phase III is actually stable.

In the large tolerance E regime, the RS solution is stable across all phases, even in phase III_{LT}. Simple calculation based on eq. (78) yields $\hat{\sigma} - \alpha = d - \epsilon < 0$, where the last inequality is proved at the end of sec. 4.2.3.

To summarize, the RS ansatz is stable in all phases, but only marginally in phases III and III_{ME}. This result may be related to the ‘simple’ structure of the version space. Consider the case of zero tolerance, $E = 0$, and two distribution vectors, \mathbf{p}_1 and \mathbf{p}_2 , in the version space. Any linear combination of these two vectors, $\mathbf{p}_t = t\mathbf{p}_1 + (1-t)\mathbf{p}_2$ with $t \in [0, 1]$, is a normalized distribution and lies in the version space. Hence the version space is convex and connected. The instability of RS ansatz, which is usually related to the appearance of many disconnected and far apart components, may therefore not take place.

5 Numerical simulations

We now present a numerical procedure to sample the space of distributions with the measure ρ . Due to the exponential growth of the version space with N this procedure is applied in practice to small values of $N \leq 10$. However the results confirm the analytical calculations reported above, and provide insights about finite-size effects. We restrict to the case of zero tolerance ($E = 0$) throughout this section.

5.1 Sampling algorithm

We resort to a Monte Carlo (MC) sampling method, in which the distribution \mathbf{p} is updated at discrete time steps. Each step corresponds to a random change of the current probability vector from \mathbf{p} to $\mathbf{p}' = \mathbf{p} + \Delta\mathbf{p}$. The move $\Delta\mathbf{p}$ must satisfy the following conditions:

Orthogonality to observable vectors: $\Delta\mathbf{p} \cdot \mathbf{v}^\mu = 0, \forall \mu$.

Normalization: $\sum_s \Delta p_s = \Delta\mathbf{p} \cdot \mathbf{1} = 0$ where $\mathbf{1} = (1, 1, \dots, 1)$.

Positivity: $p_s + \Delta p_s \geq 0, \forall s$.

The orthogonality condition ensures that the constraints keep being satisfied at all steps provided they are fulfilled by the initial value of the distribution. The same statement applies to the normalization condition; we will specify below how the initial condition is chosen. The orthogonality and normalization conditions restrict the possible directions \mathbf{w} (normalized to unity) of the move. Once this direction \mathbf{w} is chosen we determine the range $[x_{\min}; x_{\max}]$ of the allowed amplitudes x of the move $\Delta\mathbf{p} = x \mathbf{w}$ to fulfill the positivity constraint:

$$x_{\min} = \min_s \max \left(\frac{-p_s}{w_s}, \frac{1-p_s}{w_s} \right), \quad x_{\max} = \max_s \min \left(\frac{-p_s}{w_s}, \frac{1-p_s}{w_s} \right). \quad (113)$$

Any intermediate values between these two bounds may be chosen uniformly at random. Next, we calculate the entropy difference $\Delta S = S(\mathbf{p}') - S(\mathbf{p})$ and accept the move $\mathbf{p} \rightarrow \mathbf{p}'$ according to the Metropolis rule, *i.e.* with probability

$$p_{\text{accept}} = \min(1, e^{\Gamma \Delta S}), \quad (114)$$

and reject it (leave \mathbf{p} unchanged) with probability $1 - p_{\text{accept}}$. A picture illustrating one Monte Carlo step is shown in Fig. 8.




Fig. 8 Schematic picture of the version space and of one Monte Carlo step. The version space is convex, and includes the target distribution $\hat{\mathbf{p}}$. From a distribution \mathbf{p} , a direction \mathbf{w} inside the solution space is randomly chosen. We choose a point on the segment along this direction (dashed line) uniformly at random, and \mathbf{p} is updated to the chosen point \mathbf{p}' with probability p_{accept} , see eq. (114).

To implement the algorithm, it is convenient not to assume that the observables are Gaussianly distributed, as in the analytical treatment. Instead we consider the Fourier modes $\mathbf{w}^{\mathbf{k}}$ on the N -dimensional hypercube, where $\mathbf{k} = (k_1, k_2, \dots, k_N)$, with $k_i = 0, 1$ for each i , denote the wave-number configurations. The 2^N components of those Fourier modes are given by

$$(\mathbf{w}^{\mathbf{k}})_s = \frac{1}{\sqrt{2^N}} \prod_{i=1}^N (s_i)^{k_i}. \quad (115)$$

The Fourier modes $\{\mathbf{w}^{\mathbf{k}}\}_{\mathbf{k}}$ form a complete orthonormal basis of the 2^N -dimensional vector space. Note that $\mathbf{w}^{\mathbf{0}} = \mathbf{1}/\sqrt{2^N}$, where $\mathbf{0}$ is the all-zero wave-number configuration.

In the analytical calculation of sec. 4 we chose the observable \mathbf{v} to be random vectors, with Gaussian components. Any two randomly chosen vectors have very small scalar product with high probability in the large- N limit. In our numerical simulation we rather choose the M observables uniformly at random over the discrete set of $2^N - 1$ Fourier modes $\mathbf{w}^{\mathbf{k}}$, with $\mathbf{k} \neq \mathbf{0}$. This choice is convenient since the Fourier modes are orthogonal by construction, implying that the orthogonality condition mentioned above is automatically satisfied as soon as we choose the direction of the move \mathbf{w} as one of the Fourier modes outside the set of observables. The normalization condition is easily satisfied as long as we exclude $\mathbf{w}^{\mathbf{0}}$ from the set of possible directions for the move. We are thus left with a set of $2^N - M - 1$ Fourier modes, each of which is a possible direction for the Monte Carlo move. Clearly our Monte Carlo Markov Chain satisfies detailed balance and is irreducible.

5.2 Results

Using the above sampling procedure we calculate several quantities of interest including order parameters, histograms of p_s , and spin-spin correlations. The target distribution we consider is again the ISM (6) with $H = 0.5$.

5.2.1 Check of equilibration

We have run simulations for ten different samples ($N_{\text{sample}} = 10$) for sizes $N = 4, 6$ and 8 , and for two samples ($N_{\text{sample}} = 2$) for size $N = 10$ to compute the values of the order parameters. The error bars are estimated through the standard deviation σ_{sample} across the samples:

$$\text{Error bar} = \frac{\sigma_{\text{sample}}}{\sqrt{N_{\text{sample}} - 1}}. \quad (116)$$

A Monte Carlo step is defined as one attempted move, irrespective of its acceptance. Fig. 9 shows the plot of the order parameter q against the number of Monte Carlo steps, indicating how the system approaches equilibrium. From the figure we see that thermalization becomes drastically faster for larger

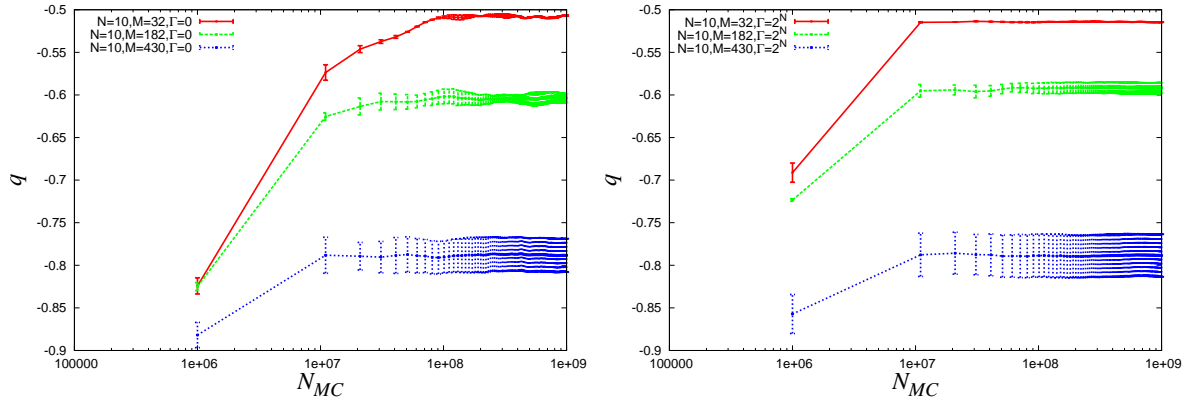


Fig. 9 Plots of q vs. the number of Monte Carlo (MC) steps for $N = 10$ spins and for different values of the number M of constraints. The left panel corresponds to $\Gamma = 0$ (no entropic bias), and the right panel to $\Gamma = 2^N = 1024$.

Γ . For example, simulations with $M = 32$ constraints and $N = 10$ spins require at least $N_{\text{MC}} \approx 2 \times 10^8$ steps for thermalization when $\Gamma = 0$, while $N_{\text{MC}} = 2 \times 10^7$ steps seem to be sufficient when $\Gamma \geq 2^{10}$. This trend is easy to understand, as an increase in the entropic bias concentrates more and more the measure around ME distribution, hence shrinking the space to be sampled.

The actual number of Monte Carlo steps used in our simulation varied depending on the quantities we wanted to estimate. To compute the values of the order parameters we chose $N_{\text{MC}} = 1 \times 10^7$ for $N = 4$ and 6 , $N_{\text{MC}} = 3 \times 10^7$ for $N = 8$, and $N_{\text{MC}} = 3 \times 10^8$ for $N = 10$. One third of those steps are discarded in the computation of the averages. To obtain accurate histogram of the distributions of the order parameters, however, we chose typically 5 times more MC steps.

5.2.2 Order parameters, marginal measures, and spin correlations

To confirm our analytical prediction and to estimate finite-size effects we compute the order parameters for various values of M, Γ and N , and report the results in Fig. 10. Those figures clearly show that the analytical prediction, black curves, are consistent with the numerical results for sizes as small as $N = 10$. The phase transitions are well captured. In particular, the phase transition to the ME phase is clearly seen in the behaviour of d (right lower panel). This agreement show that our analytical findings, derived in the infinite- N limit, are numerically accurate even for small sizes.

We have also computed the histograms of the single-configuration probabilities p_s , with the results shown in Figs. 11 and 12. Due to the symmetry of the ISM the target probabilities p_s depend only on the number of, say, -1 spins in the configuration s , which we call N_s^- . We checked that this symmetry is approximately recovered in the histograms despite the noise introduced by the Monte Carlo sampling. Therefore we show only one among the p_s with the same N_s^- in those figures.

Fig. 11 shows the histograms corresponding to different N_s , for $N = 8$ and $\Gamma = 0$, and two values of the number M of constraints. Two characteristic shapes of the histograms emerge. One looks like roughly

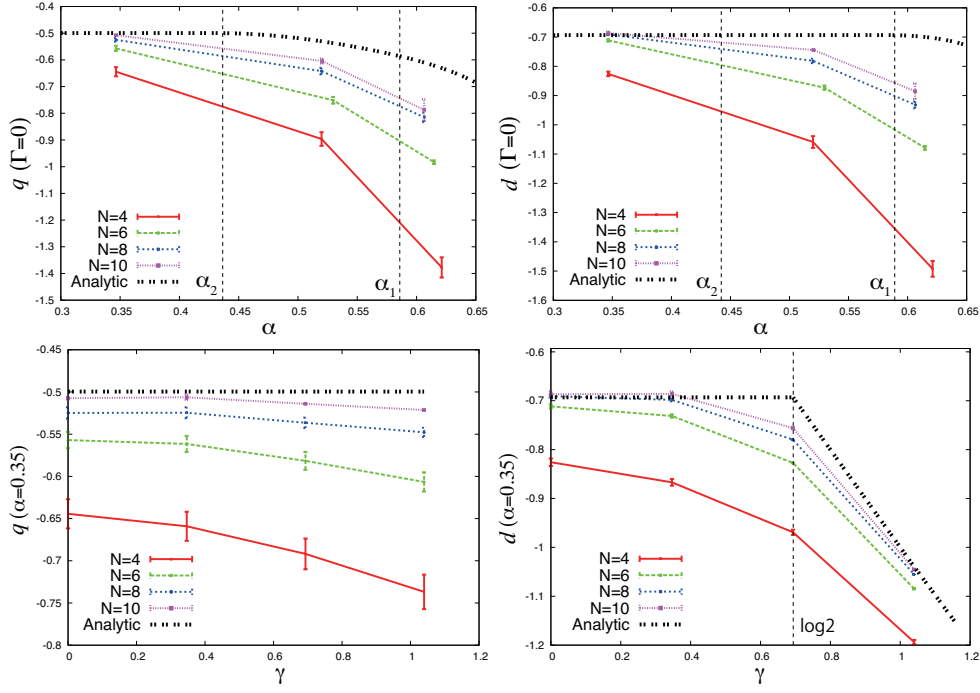


Fig. 10 Order parameters q (left) and d (right) of the ISM with $H = 0.5$, computed from Monte Carlo simulations and plotted vs. α (upper row) and γ (lower row). Analytical predictions are shown with the black curves. The phase transition from small to large Γ can be best guessed in the lower, right panel. Finite-size effects seem to be stronger for larger α .




Fig. 11 Histograms of single-configuration probabilities p_s for $N = 8$ and $\Gamma = 0$; the numbers N_s^- of negative spins in the configurations are indicated on the figure. The left panel corresponds to $M = 30$, and the right panel to $M = 80$.

Gaussian, while the other typical behaviour resembles an exponential distribution. These shapes are in very good agreement with the two possible functional dependences of $\rho_s(p_s)$ upon p_s given in Fig. 3. A larger number of Gaussian-like histograms is found for $M = 80$ than for $M = 30$, implying that more target probabilities are learned in the former case than in the latter, as expected from the theory. The size and sample dependence of the histograms of p_0 (corresponding to the configuration **1** with $N_s^- = 0$ minus spins) are examined in Fig. 12. The corresponding target probability \hat{p}_1 is the largest one of the ISM and is accurately learned for the values of M corresponding to the figures. From the left panel of Fig. 12, we see that our analytical prediction regarding $\rho_s(p_s)$ being centered in $\hat{p}_s - 1/M$, becomes more and more accurate as the system size increases. An excellent agreement with the prediction is reached for $N = 10$. Sample-to-sample fluctuations of the histogram of p_1 are shown in the right panel. The choice




Fig. 12 (Left) System-size dependence of histograms of p_0 (corresponding to the configuration **1** all plus spins), for fixed $\alpha \approx 0.598$ and $\Gamma = 0$. For each histogram, the rightmost dashed line show the location of the target value \hat{p}_1 , while the leftmost dashed line shows $\hat{p}_1 - 1/M$. (Right) Sample-to-sample fluctuations of the histogram of p_0 for $N = 8$, $M = 80$ and $\Gamma = 0$. The dashed vertical line represents the target value.

of the constraints produces moderate fluctuations in the location of the peak height of the Gaussian-like histograms.

Last of all we show in Fig. 13 the histograms of multi-spin correlations, for specific subsets of spins. We consider in particular

$$c_{12} = \sum_s p_s s_1 s_2, \quad c_{123} = \sum_s p_s s_1 s_2 s_3. \quad (117)$$

We see that the predicted values of c are in fair agreement with the target distribution values, but both




Fig. 13 Histograms of pairwise (c_{12} , left) and triplet (c_{123} , right) correlations for a system of $N = 8$ spins. Vertical lines indicate the target distribution values. Each histogram corresponds to one sample of $M = 80$ observables.

the thermal and the sample-to-sample fluctuations are rather large, compared to the results of Fig. 12.

6 Conclusion

In this paper we have investigated the properties of the space of probability distributions (over a large set of discrete variable configurations), constrained to reproduce many average values of observables,

computed according to a target distribution $\hat{\mathbf{p}}$. For zero tolerance $E = 0$, the space of admissible distributions \mathbf{p} defines the version space. The version space contains many distributions of interest, in addition to the target distribution $\hat{\mathbf{p}}$, such as the center-of-mass distribution, $\langle \mathbf{p} \rangle$, which is the flat average over all admissible distributions, and the maximum entropy distribution, \mathbf{p}_{ME} . In the case of finite tolerance $E > 0$ the version space acquires a probabilistic meaning. We introduce a probability measure ρ over the space of distributions, giving more weight to the distributions \mathbf{p} in good agreement with the target values of the observables. The measure ρ may furthermore be biased to favour distributions \mathbf{p} with large entropies $S(\mathbf{p})$. To do so we introduce a multiplicative factor in the measure ρ , growing exponentially with $\Gamma S(\mathbf{p})$. The coefficient Γ acts as an inverse temperature: $\Gamma = 0$ allows us to find back the unbiased measure, while, for $\Gamma \rightarrow \infty$, the measure becomes fully concentrated around \mathbf{p}_{ME} . Varying the value of Γ allows us to understand how effective is the maximum entropy principle to approximate the target distribution $\hat{\mathbf{p}}$.

To compute analytically the volume of the version space and its main properties we have assumed that observables were quenched random variables, and we have ignored correlations between different values of the observable in different variable configurations. This assumption is not realistic, as physical observables are generally rather smooth functions of the configurations. As a result of this simplifying assumption, we have been able to compute the typical distances between the target, the centre-of-mass, and the maximum entropy distributions, as well as the typical fluctuations around the centre of mass. The calculation was done within the replica symmetric framework, and we have checked that our solution was locally stable against replica-symmetry-breaking fluctuations (replicon modes). A major outcome of our calculation is the notion of learning edge, which separates learned from as-yet-unknown configuration probabilities. Learned probabilities correspond to large target probabilities, while the probabilities of configurations associated to small target probabilities remain unknown. As the number of observables increases the learning edge moves to smaller and smaller probability values, implying that learning proceeds. The decrease of the learning edge is not generally accompanied by a decrease of the distance between the centre-of-mass and the target distribution, unless it hits the probabilities contributing most to that distance at specific critical points.

Numerical simulations were performed to confirm our asymptotic and analytical predictions, and to quantify finite-size effects. Due to the exponential growth of the dimension of the version space with the number N of variables we are limited to very small sizes, in practice $N \leq 10$. Nevertheless results are in good agreement with our analytical results and finite-size effects do not seem to alter our asymptotic results in a significant way. In particular our predictions for the existence of a learning edge and for the functional forms of the marginal measures over the probabilities of configurations are remarkably confirmed by the numerics.

We have also studied the role of the entropic bias for a given number of measured observables. Our calculation shows that the maximum entropy distribution is not closer to the target distribution than any other randomly chosen distribution in the version space. This negative result is due to our choice of fully uncorrelated observables. It is quite likely that introducing an entropic bias becomes efficient and improves learning performances if we require that both the observables and the target distribution satisfy some smoothness criteria. Some preliminary results, obtained without assuming that the observable values are uncorrelated from configuration to configuration, are reported in [9] and seem to support this guess. Unfortunately the analytical calculations with ‘correlated’ observables are very involved, and we have not been able to make substantial progresses so far. The theoretical importance and the practical relevance of the issues addressed here, such as how the number of constraints should depend on the smoothness of the observables and of the target distribution, or whether the phase transitions at distinct steps of the learning process found in the random uncorrelated case studied here exist in the correlated case, are strong incentives to extend this study to realistic observable ensembles. Another valuable direction for further research would be to search for ‘good’ sets of observables. While we have restricted here to the case of M independently drawn observables, it would be interesting to optimize the choice of those observables to make the volume of the version space as small as possible, and to get as close as possible to the target distribution. In this regard, we have extended the present analysis to the case of a finite number of replicas, $n \neq 0$, the result of which will be reported soon.

Acknowledgments

The authors are grateful to U. Ferrari, Y. Kabashima, J. Lebowitz, and T. Mora for fruitful discussions. T. O. acknowledges the support by Grant-in-Aid for JSPS Fellows, as well as the JSPS Core-to-Core Program “Non-equilibrium dynamics of soft matter and information”. S. C. and R. M. acknowledge financial support from the [EU-]FP7 FET OPEN project Enlightenment 284801 and Agence Nationale de la Recherche Coevstat project (ANR-13-BS04-0012-01). Numerical calculations were partly carried out on the TSUBAME2.5 supercomputer in the Tokyo Institute of Technology.

A Lesson from no-constraint case for finite Γ

If Γ is finite in the no-constraint case, the volume is written as

$$V(M=0, \Gamma) = \int_{-\infty}^{\infty} d\Lambda e^{\Lambda} \int_0^{\infty} \prod_s \left(dp_s e^{-\Lambda p_s - \Gamma p_s \log p_s} \right). \quad (118)$$

Now, the saddle point with respect to p_s is quite simple

$$p_s^* = \tilde{p} = e^{-1-\Lambda/\Gamma}. \quad (119)$$

This is the origin of the solution assumed in sec. 4.3.1. Assuming the saddle-point approximation is correct, we can write the integration as

$$\int dp_s e^{-\Lambda p_s - \Gamma p_s \log p_s} = \sqrt{2\pi} \sqrt{\frac{e^{-1-\frac{\Lambda}{\Gamma}}}{\Gamma}} e^{\Gamma e^{-1-\frac{\Lambda}{\Gamma}}} \quad (120)$$

and the log volume is

$$F = \Lambda + 2^N \left\{ \Gamma e^{-1-\frac{\Lambda}{\Gamma}} + \frac{1}{2} \left(\log(2\pi) - 1 - \frac{\Lambda}{\Gamma} - \log \Gamma \right) \right\}. \quad (121)$$

Taking a variation with respect to Λ , we get

$$\frac{1}{2^N} = e^{-1-\frac{\Lambda}{\Gamma}} + \frac{1}{2\Gamma}. \quad (122)$$

We know $\Lambda = 2^N$ at $\Gamma = 0$. Thus, it is natural to think that the saddle point $\tilde{p} = e^{-1-\Lambda/\Gamma}$ takes meaningful value only for $\Gamma \geq 2^N$, and otherwise \tilde{p} rapidly becomes zero and the result comes back to the case $\Gamma = 0$. Hence, eq. (122) is satisfied in the leading scale as $2^{-N} \doteq e^{-1-\Lambda/\Gamma}$. Hence, $\Lambda = (N \log 2 - 1)\Gamma$ and the term $1/2\Gamma$ becomes subleading one. Substituting this, we get

$$F \approx (N \log 2)\Gamma - 2^N \left(N \log 2 + \log \Gamma + \frac{1}{2}(1 - \log 2\pi) \right). \quad (123)$$

The leading scale of the log volume is thus dominated by Γ .

B Stability analysis

The replica generating function is given by

$$\phi(n) = \frac{1}{2} \sum_{a \leq b} Q_{ab} Q'_{ab} - \frac{M}{2} \text{Tr} \log (E + Q) + \frac{M}{2} n \log E + \sum_s \log \Theta_s, \quad (124)$$

Here Tr denotes the trace of matrix. We write the exponent in Θ_s in eq. (39) as

$$f_s(\{p_s^a\}_a) = - \sum_a \Lambda_a (p_s^a - \hat{p}_s) - \frac{1}{2} \sum_{a \leq b} Q'_{ab} (p_s^a - \hat{p}_s)(p_s^b - \hat{p}_s) - \Gamma \sum_a p_s^a \log p_s^a. \quad (125)$$

Let us consider some small fluctuations around the RS saddle-point order parameters:

$$Q_{ab} = Q_{ab}^{\text{RS}} + x_{ab}, \quad Q'_{ab} = Q_{ab}^{\text{RS}} + 2\hat{x}_{ab}. \quad (126)$$

Then,

$$\frac{1}{2} \sum_{a \leq b} Q_{ab} Q'_{ab} \approx \frac{1}{2} \sum_{a \leq b} Q_{ab}^{\text{RS}} Q_{ab}^{\text{RS}} + \sum_{a \leq b} x_{ab} \hat{x}_{ab}, \quad (127)$$

$$\text{Tr} \log (E + Q) \approx \text{Tr} \log (E + Q^{\text{RS}}) - \frac{1}{2} \text{Tr} (Ax)^2. \quad (128)$$

where we define $A = (E + Q^{\text{RS}})^{-1}$, and

$$f_s = f_s^{\text{RS}} - \sum_{a \leq b} \hat{x}_{ab} \bar{p}_s^a \bar{p}_s^b = f_s^{\text{RS}} + \Delta f_s, \quad (129)$$

with $\bar{p}_s^a = p_s^a - \hat{p}_s$. Thus

$$\log \Theta_s \approx \log \text{Tr} e^{f_s^{\text{RS}}} \left(1 + \Delta f_s + \frac{1}{2} \Delta f_s^2 \right) \approx \log \Theta^{\text{RS}} + \frac{1}{2} \langle \Delta f_s^2 \rangle - \frac{1}{2} \langle \Delta f_s \rangle_s^2, \quad (130)$$

where we have introduced the notation

$$\langle \cdots \rangle_s = \frac{1}{\Theta_s^{\text{RS}}} \text{Tr} e^{f_s^{\text{RS}}} (\cdots). \quad (131)$$

Note that we have omitted all the first-order terms, which vanish at the saddle point. Thus we get $\phi(n) = \phi^{\text{RS}} + \Delta$ with

$$2\Delta = 2 \sum_{a \leq b} x_{ab} \hat{x}_{ab} + \frac{M}{2} \text{Tr}(Ax)^2 + \sum_{a \leq b} \sum_{c \leq d} \hat{x}_{ab} \hat{x}_{cd} \sum_s \left(\langle \bar{p}_s^a \bar{p}_s^b \bar{p}_s^c \bar{p}_s^d \rangle_s - \langle \bar{p}_s^a \bar{p}_s^b \rangle_s \langle \bar{p}_s^c \bar{p}_s^d \rangle_s \right). \quad (132)$$

We write $A_{aa} \equiv X$ and $A_{ab} \equiv Y$ ($a \neq b$), and those components are easily calculated

$$X = \frac{E + Q + (n-2)R}{(E + Q + (n-1)R)(E + Q - R)}, \quad Y = -\frac{R}{(E + Q + (n-1)R)(E + Q - R)}, \quad (133)$$

$$X + (n-1)Y = \frac{1}{(E + Q + (n-1)R)}, \quad X - Y = \frac{1}{(E + Q - R)}. \quad (134)$$

Eq. (132) is a quadratic form with respect to $\{x_{ab}\}$ and $\{\hat{x}_{ab}\}$ and can be written as $2\Delta = \mathbf{V}^t G \mathbf{V}$, where \mathbf{V} is a column vector with components $\{x_{ab}\}$ and $\{\hat{x}_{ab}\}$, and \mathbf{V}^t is its transpose. We now want to determine whether the Hessian matrix G has unstable modes. The most likely candidate lies in the replicon eigenspace, which is spanned by the vectors whose components (x_{ab}, \hat{x}_{ab}) depend only on whether their replica indices are equal or different to two fixed values $a = \theta$ and $b = \eta$ [8]. In the replicon space, diagonal ($a = b$) fluctuations tend to be irrelevant since other (transverse, longitudinal) modes span the n -dimensional space (x_{aa}, \hat{x}_{aa}) . Hence we set hereafter $x_{aa} = \hat{x}_{aa} = 0$.

We arrange the components of \mathbf{V} by ordering x as

$$\mathbf{v} = (x_{12}, x_{13}, \cdots, x_{1n}, x_{23}, \cdots, x_{n-1,n}), \quad (135)$$

and \hat{x} as well. The Hessian G has the following form

$$G = \left(\begin{array}{ccc|ccc} S & T \cdots T & U \cdots U & & & \\ & \ddots & & & I & \\ & & \ddots & & & \\ T \cdots T & U \cdots U & S & & & \\ & & & \hat{S} & \hat{T} \cdots \hat{T} & \hat{U} \cdots \hat{U} \\ & I & & & \ddots & \\ & & & \hat{T} \cdots \hat{T} & \hat{U} \cdots \hat{U} & \hat{S} \end{array} \right), \quad (136)$$

where I is the identity matrix, and

$$\hat{S} = \sum_s \left(\langle (\bar{p}_s^a \bar{p}_s^b)^2 \rangle_s - \langle \bar{p}_s^a \bar{p}_s^b \rangle_s^2 \right), \quad (137)$$

$$\hat{T} = \sum_s \left(\langle (\bar{p}_s^a)^2 \bar{p}_s^b \bar{p}_s^c \rangle_s - \langle \bar{p}_s^a \bar{p}_s^b \rangle_s \langle \bar{p}_s^a \bar{p}_s^c \rangle_s \right), \quad (138)$$

$$\hat{U} = \sum_s \left(\langle \bar{p}_s^a \bar{p}_s^b \bar{p}_s^c \bar{p}_s^d \rangle_s - \langle \bar{p}_s^a \bar{p}_s^b \rangle_s \langle \bar{p}_s^c \bar{p}_s^d \rangle_s \right), \quad (139)$$

$$S = M(X^2 + Y^2), \quad T = MY(X + Y), \quad U = 2MY^2. \quad (140)$$

Let us find the eigenvectors of this matrix. The first eigenvector \mathbf{V}_1 is obtained by assuming $x_{ab} = a$ and $\hat{x}_{ab} = \hat{a}$ for any λ, ν . The upper half (in the matrix drawn in eq. (136)) of the eigenvalue equation $G\mathbf{V}_1 = \lambda_1 \mathbf{V}_1$ gives

$$\left(S + 2(n-2)T + \frac{1}{2}(n-2)(n-3)U \right) a + \hat{a} = \lambda_1 a, \quad (141)$$

and the lower half yields

$$a + \left(\hat{S} + 2(n-2)\hat{T} + \frac{1}{2}(n-2)(n-3)\hat{U} \right) \hat{a} = \lambda_1 \hat{a}. \quad (142)$$

These equations have the eigenvalue λ_1 with non-vanishing a, \hat{a} , which must satisfies the following relation

$$\lambda_1^2 - (C_1 + \hat{C}_1)\lambda_1 + (C_1 \hat{C}_1 - 1) = 0, \quad (143)$$

where $C_1 = S + 2(n-2)T + (n-2)(n-3)U/2$ and $\hat{C}_1 = \hat{S} + 2(n-2)\hat{T} + (n-2)(n-3)\hat{U}/2$. This equation says that this mode spans a two-dimensional space.

The next type of solution \mathbf{V}_2 is obtained by selecting a replica index, say, $\theta = 1$. This solution \mathbf{V}_2 has $x_{ab} = b$ and $\hat{x}_{ab} = \hat{b}$ when λ or ν is equal to 1, $y^{\lambda\nu} = c$ and $\hat{y}^{\lambda\nu} = \hat{c}$ otherwise. The first row of the eigenvalue equation $G\mathbf{V}_2 = \lambda_2\mathbf{V}_2$ gives

$$(S + (n-2)T)b + \left((n-2)T + \frac{1}{2}(n-2)(n-3)U\right)c + \hat{b} = \lambda_2 b, \quad (144)$$

and the first row of the lower half of matrix G in eq. (136) yields

$$b + \left(\hat{S} + (n-2)\hat{T}\right)\hat{b} + \left((n-2)\hat{T} + \frac{1}{2}(n-2)(n-3)\hat{U}\right)\hat{c} = \lambda_2 \hat{b}. \quad (145)$$

We now impose the orthogonality condition over \mathbf{V}_2 and \mathbf{V}_1 . This leads to

$$(n-1)b + \frac{1}{2}(n-1)(n-2)c = 0, \quad (146)$$

and the same relation holds for \hat{b} and \hat{c} . Substituting these conditions, we get

$$(S + (n-4)T - (n-3)U)b + \hat{b} = \lambda_2 b, \quad \left(\hat{S} + (n-4)\hat{T} - (n-3)\hat{U}\right)\hat{b} + b = \lambda_2 \hat{b}, \quad (147)$$

which leads to

$$\lambda_2^2 - (C_2 + \hat{C}_2)\lambda_2 + (C_2\hat{C}_2 - 1) = 0, \quad (148)$$

where $C_2 = S + (n-4)T - (n-3)U$ and $\hat{C}_2 = \hat{S} + (n-4)\hat{T} - (n-3)\hat{U}$. As there are n possible choices for the replica index θ and two eigenvalues/eigenvectors \mathbf{V}_1 and \mathbf{V}_2 for each choice, this particular subspace is of dimension $2n$.

The third mode \mathbf{V}_3 is obtained by treating two replicas θ, ω as special ones. This solution \mathbf{V}_3 has $x_{\theta\omega} = d$ and $\hat{x}_{\theta\omega} = \hat{d}$, $x_{\theta a} = e$ and $\hat{x}_{\omega a} = \hat{e}$, and $x_{ab} = f$ and $\hat{x}_{ab} = \hat{f}$ otherwise. The orthogonality condition with \mathbf{V}_2 is

$$b(d + (n-2)e) + c\left((n-2)e + \frac{1}{2}(n-2)(n-5)f\right) = b(d + (n-4)e - (n-5)f) = 0, \quad (149)$$

where we use eq. (146). The one with \mathbf{V}_1 is

$$d + 2(n-2)e + \frac{1}{2}(n-2)(n-5)f = 0. \quad (150)$$

These relations mean

$$e = -\frac{d}{n-2}, \quad \frac{1}{2}(n-2)(n-5)f = -d - 2(n-2)e = d. \quad (151)$$

Similar relations hold for the hat variables. According to these relations, the first rows of the upper and lower halves of the eigenvalue equation $G\mathbf{V}_3 = \lambda_3\mathbf{V}_3$ yield

$$dS + 2(n-2)eT + \frac{1}{2}(n-2)(n-5)fU + \hat{d} = (S - 2T + U)d + \hat{d} = \lambda_3 d, \quad (152)$$

$$d + (\hat{S} - 2\hat{T} + \hat{U})\hat{d} = \lambda_3 \hat{d}, \quad (153)$$

Thus we obtain

$$\lambda_3^2 - (C_3 + \hat{C}_3)\lambda_3 + (C_3\hat{C}_3 - 1) = 0, \quad (154)$$

where $C_3 = S - 2T + U$ and $\hat{C}_3 = \hat{S} - 2\hat{T} + \hat{U}$. This solution spans a $n(n-3)$ -dimensional space, implying that no eigenmode are left.

For the stability of the saddle point, all of the eigenvalues must be non-negative. This condition corresponds to

$$\forall i, \quad C_i \hat{C}_i \leq 1. \quad (155)$$

This condition takes into account the fact that Q'_{ab} is originally a pure imaginary variable, which means that $\delta Q_{ab}\delta Q'_{ab}$ is associated with a multiplicative factor i and $\delta Q'_{ab}\delta Q_{ab}$ acquires a factor -1 . Hence, if we change variable from Q'_{ab} to iQ'_{ab} , the diagonal block in the lower half part of G gets a factor -1 , and the off-diagonal part becomes iI , which leads to the positivity condition in eq. (155).

The replicon mode corresponds to λ_3 and \mathbf{V}_3 . Thus, the stability condition is

$$(S - 2T + U)(\hat{S} - 2\hat{T} + \hat{U}) \leq 1. \quad (156)$$

A straightforward calculation shows that

$$S - 2T + U = M(X - Y)^2 = \frac{M}{(E + Q - R)^2}. \quad (157)$$

We now turn to the calculation of $\hat{S} - 2\hat{T} + \hat{U}$. Within the RS assumption the average $\langle O \rangle_s$, where $O = O_a O_b \cdots O_c$, can be expressed as

$$\langle O \rangle_s = \frac{1}{\int Dz X_s^n} \int Dz X_s^n \langle O_a \rangle_{X_s} \cdots \langle O_c \rangle_{X_s}, \quad (158)$$

where X_s was defined in eq. (47). Thus,

$$\begin{aligned} \hat{S} - 2\hat{T} + \hat{U} &= \sum_s \left\{ \left\langle (\bar{p}_s^a)^2 (\bar{p}_s^b)^2 \right\rangle_s - 2 \left\langle (\bar{p}_s^a)^2 \bar{p}_s^b \bar{p}_s^c \right\rangle_s - \left\langle \bar{p}_s^a \bar{p}_s^b \bar{p}_s^c \bar{p}_s^d \right\rangle_s \right\} \\ &= \sum_s \frac{1}{\int Dz X_s^n} \int Dz X_s^n \left(\langle \bar{p}_s^2 \rangle_{X_s} - \langle \bar{p}_s \rangle_{X_s}^2 \right)^2. \end{aligned} \quad (159)$$

Noticing that

$$\langle \bar{p}_s^2 \rangle_{X_s} - \langle \bar{p}_s \rangle_{X_s}^2 = \frac{\partial^2 \log X_s}{\partial \Lambda^2}, \quad (160)$$

the stability condition may be written, in the $n \rightarrow 0$ limit, under the form shown in eq. (111).

References

1. Jaynes, E.T.: Information Theory and Statistical Mechanics, *Physical Review* **106**, 620-630 (1957); Jaynes, E.T.: Information Theory and Statistical Mechanics II, *Physical Review* **108**, 171-190 (1957).
2. Balian, R.: Statistical mechanics and the maximum entropy method, in From Statistical Physics to Statistical Inference and Back, Grassberger, P., Nadal, J.P. (eds), NATO ASI series 428, Kluwer Academics Publisher, p. 11-43 (1994).
3. Cover, T.M., Thomas, J.A.: Elements of information theory, Wiley Series in Telecommunications and Signal Processing, Wiley (1991).
4. Jaynes, E.T.: On the Rationale of Maximum Entropy Methods, *Proc. IEEE* **70**, 939-952 (1982).
5. Jaynes, E.T.: Monkeys, kangaroos and N , in Maximum-Entropy and Bayesian Methods in Applied Statistics, J.H. Justice (ed.), Cambridge University Press, Cambridge, p. 26-58 (1986).
6. Tikochinsky, Y., Tishby, N.Z., Levine, R.D.: Alternative approach to maximum-entropy inference, *Phys. Rev. A* **30**, 2638-2644 (1984).
7. Bialek, W.: Biophysics: Searching for Principles, Princeton University Press, Princeton (2012).
8. De Almeida, J.R.L., Thouless, D.J.: Stability of the Sherrington-Kirkpatrick solution of a spin glass model, *Journal of Physics A: Mathematical and General* **11**, 983-990 (1978).
9. Obuchi, T., Monasson, R.: Submitted to J. Phys. Conf. Ser. (2015).